Adaptive eLearning for Grid Computing

by

Kathryn Cassidy

Bachelor of Science

A Thesis submitted to
The University of Dublin
for the degree of

Doctor of Philosophy

Department of Computer Science,
University of Dublin,
Trinity College

March, 2012

Declaration

I declare that this thesis has not been submitted as an exercise for a degree at this or any other university and it is entirely my own work.
I agree to deposit this thesis in the University's open access institutional repository or allow the library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement

Signature of Author ....

Kathryn Cassidy, 23th March, 2012

Summary

Adaptive eLearning appears well suited to Grid education because of the distributed and heterogeneous nature of grid users and their asynchronous training demand. This thesis explores the application of adaptive eLearning techniques to grid computing education. It looks at integration of an adaptive eLearning application, and associated courses, with a training infrastructure which provides a fair replica of a production Grid.

It explores the potential for extending this integration to the practical components of online courses, allowing learners to access practical environments within the eLearning application and enabling the adaptive eLearning tool to capture relevant information about learners' conduct of practical exercises within this environment, so-called *"XeL"* for eXecutable eLearning, the core hypothesis of this thesis. Captured information can then be used by the adaptive eLearning application for further course personalisation, potentially enhancing the learning experience.

A prototype implementation of an infrastructure and tools to explore these ideas is described. The prototype validated that the proposed infrastructure is feasible.

Several experiments were conducted to gather data about the efficacy or otherwise of the approach and the results of these experiments are presented and some tentative conclusions drawn as to the benefits of the XeL approach. Learners found XeL both intuitive and useful, and there is some evidence for increased learning, and enhanced suitability for remote learners when using XeL.

Future areas of research are proposed in order to further elaborate on the ideas presented here.

Acknowledgements

I extend my sincere thanks to my supervisor Dr.Brian Coghlan. His broad knowledge, endless patience and thorough attention to detail have been of enormous help to me over the past five years. His ideas and insight were key in expanding my initial ideas about this research. I am also indebted to the HEA PRTLI3 CosmoGrid project and the HEA PRTLI4 e-INIS project (both led by Prof.Luke Drury, DIAS), whose generous support made this research possible.

The members past and present of the Computer Architecture and Grid Research Group in Trinity College Dublin acted as mentors and sounding-boards for many of the ideas in this thesis. Working with John Walsh on the *Introduction to Developing Grid Applications* course laid the foundations for this research. He, along with Dr.David O'Callaghan, Dr.Eamonn Kenny, Dr.Stuart Kenny, Dr.Stephen Childs, Dr.Geoff Quigley and Dr.John Ryan kindly assisted me with aspects of the eLGrid infrastructure. I must also thank Dr.Ronan Watson and Peter Lavin, my team-mates on the Statistics course, for their assistance with the R statistical package and their agreement to use a subset of my data for our group project, a useful dry-run for my own analysis. Indeed, to all of the CAG, KDEG and SCSS members who gave their time to help with emergency testing I am grateful.

Thanks are also due to KDEG for allowing me to use APeLS and ACCT within eLGrid, and special mention must go to Dr.Declan Dagger who provided guidance on usage of these tools and much helpful feedback on aspects of this research.

The organisers and attendees of the *PIC Simulations* workshop and the *EELA Grid School* in TCD in 2008 were trusting enough to use eLGrid and several eLGrid courses as a central part of their training. Similarly, the students of courses 4BA9 (2009) and CS4021 (2009, 2010) in the Computer Science Degree programme of Trinity College Dublin all willingly subjected themselves to my tortuous tests and surveys.

I have been involved, over the past five years, in several Grid projects that have significantly enhanced this research. I would like to specifically thank the members of the ICEAGE Consortium for kindly allowing me to tag along. Also, the organisers of the ISSGC, which I attended both in the capacity of student and organiser. The members of the OGF Education and Training Community Group have all contributed in one way or another to this research. Membership of this group gave me the opportunity to discuss my work with experts in the grid education field for which I am truly grateful. I am also indebted to OGF themselves, who provided me with a scholarship to attend GGF17 in Tokyo, Japan in May of 2006.

I would like to thank my parents and family for their continued assistance and unwavering support throughout my "scenic-route" through education. Thanks also to all of the kenshi of the TCD Kendo club, the Dublin Kendo club, Kendo na hÉireann, and the Irish National Kendo Team, without whose inspiration, support and friendship, this thesis might have been completed a good deal earlier.

Finally to Martin, thank you... I tak twoją zawsze będę, i przy tobie zawsze siędę...

To these and to everyone else who helped on the way, I say a heartfelt "thank you".

Contents

List of Figures

List of Tables

Glossary

4BA9	Advanced Computer Architecture final year course at TCD
ABC	Assess By Computer, for computer-assisted marking
AccessGrid	A virtual collaboration tool
ACCT	Adaptive Course Construction Toolkit
ACTSim	An adaptive simulator for soft skills training
Adobe Connect	A web conferencing tool
AHS	Adaptive Hypermedia System
ANOVA	Analysis of Variance statistical method
Apache	An open source web server developed by the Apache Foundation
APeLS	Adaptive Personalised eLearning Service
ARC	A grid middleware originating in Scandanavia
Assima	A simulator supporting screen-capture and replay
Blackboard	A commercial learning management system
Bonferroni adjust.	A statistical method used to address the increased probability of type I error in multiple comparisons
BugScope	Educational project providing online free interactive access to a scanning electron microscope
CA	Certificate Authority
CAG	Computer Architecture and Grid research group at TCD
Captivate	A simulator supporting screen-capture and replay

CBT	Computer-Based Training
CE	A gLite Computing Element
CERN	European Organization for Nuclear Research
Cloud	Provisioning model providing on-demand access to virtualised resources via a computer network
Condor	An opportunistic scheduling system from Univ.Wisconsin
CS4021	Advanced Computer Architecture final year course at TCD (formerly 4BA9)
CSS	Cascading Style Sheet
CVS	Concurrent Versions System - an open-source version control system
dCache	A Grid storage service originating in Germany
DCI	Distributed Computing Infrastructure
DemoGrid	An instructional grid environment
DIAS	The Dublin Institute for Advanced Studies
DN	Distinguished Name, a component of an X.509 certificate
DNS	Domain Name Service
EC	European Commission
EDG	European DataGrid project
EduGAIN	A European project providing a federated identity service for participating federations
EELA	EU project linking to the Latin American Grid
EGEE	Enabling Grids for E-SciencE
EGI	European Grid Infrastructure
EGI.eu	The EGI headquarters entity in Amsterdam
EGI-InSPIRE	An EU FP7 project to support the European Grid Infrastructure (EGI)
e-INIS	The Irish National e-Infrastructure

e-IRG	e-Infrastructure Reflection Group
eLearning	Electronic Learning
ELeGI	European Learning Grid Infrastructure
eLGrid	Grid eLearning prototype instrument constructed for this thesis
Ellipse	A learning plugin (from O'Reilly) for the Eclipse IDE
EMI	European Middleware Initiative
e-Science Talk	A Grid dissemination project (formerly GridTalk)
ETTF	Education and Training Task Force of the e-IRG
EU	European Union
EUChiaGRID	EU project linking to the Chinese Grid
EUDCISS	European Distributed Computing Infrastructures Summer School on Grid, Cloud and Desktop Computing Services
EU-IndiaGRID	EU project linking to the Indian Grid
EUMEDGRID	EU project linking to the Mediterranean Grid
Fibonacci number	The sequence F_n=F_{n-1}+F_{n-2}, seeded by F_0=0, F_1=1
FP7	Framework Program 7
F-test	An ANOVA statistical test of significance for multiple groups
GANGA	Python-based graphical user interface for Grid
GB	GigaByte
Genius	A grid-enabled portal originating from Italy
GGF	Global Grid Forum, precursor to the Open Grid Forum OGF
GILDA	Grid INFN Virtual Laboratory for Dissemination Activities
gLite	A grid middleware originating in the EU EDG/EGEE project series
Globus	A widely-used Grid toolkit

GMA	Grid Monitoring System
Grid	A computing paradigm involving shared resources which span organisational boundaries
GridBriefings	Periodic Grid newsletter maintained by e-Science Talk
GridBuilder	A tool to configure and start new VMs
GridCafe	Grid dissemination outlet maintained by e-Science Talk
GridCast	Grid dissemination outlet maintained by e-Science Talk
GridGuide	Grid dissemination outlet maintained by e-Science Talk
Grid-Ireland	National Grid Initiative for Ireland
GridKA	A summer school held by Karlsruhe Institute of Technology
GridSite	A grid-enabled webserver originating in Manchester
GridTalk	A Grid dissemination project (now e-Science Talk)
GRNET	Greek Research and Technology Network, the National Grid Initiative (and NREN) for Greece
HEA	The Irish Higher Education Authority
HEAnet	The NREN for Ireland
HEP	High-Energy Physics
HTML	HyperText Markup Language
Hypergraph	An applet for viewing tree structures in hyperbolic space
IaaS	Infrastructure as a Service
ICEAGE	International Collaboration to Extend and Advance Grid Education
IDE	Integrated Development Environment

in2p3	Institut National de Physique Nucléaire et de Physique des Particules (French National Nuclear and Particle Physics institute)
IP	Internet Protocol
IQR	Inter-Quartile Range, a statistical measure
iSGTW	International Science Grid This Week, a Grid on-line magazine
ISSGC	International Summer School of Grid Computing
IWSGC	International Winter School of Grid Computing
Java	An object-oriented language originating from Sun Microsystems
JDL	Job Description Language, originating in the European Datagrid project
JSP	Java Server Pages
KB	KiloByte
KDEG	Knowledge and Data Engineering Group at TCD
KIT	Karlsruhe Institute of Technology
LCFG	A fabric management system originating from Edinburgh
Likert Scale	A scale for deriving ordinal data from survey questionnaires
Linux	A UNIX-like open source multi-user operating system originating in Finland
LMS	Learning Management System
LVM	Logical Volume Manager
MB	MegaByte
MD	Migrating Desktop, a graphical user interface for Grid originating in Poznan
Moodle	An open-source learning management system
MPI	Message Passing Interface

MyProxy	A repository for proxy X.509 certificates
NAT	Network Address Translation
NFS	A network file system
NGI	National Grid Initiative
NREN	National Research and Education Network
NTP	Network Time Protocol
OGF	Open Grid Forum
OGF ET-CG	Open Grid Forum Education and Training Community Group
OGSA-DAI	A grid-enabled database integration service
OpenCA	Open source software for implementing a certificate authority
OpsCentre	Grid-Ireland Operations Centre at TCD
OSG	Open Science Grid
Pearson Coefficient	A statistical correlation measure
P-Grade	A grid-enabled portal originating from Hungary
PIC	Particle-in-Cell method of solving a certain class of partial differential equations
PKI	Public Key Infrastructure
PRTLI	Programme for Research in Third-Level Institutions, a series of funding programmes run by the HEA
QoS	Quality of Service
Quattor	A fabric management system originating from CERN
RA	Registration Authority
Rapidbuilder	A simulator supporting screen-capture and replay
R-GMA	Relational Grid Monitoring System
SaaS	Software as a Service
Sakai	An open-source learning management system
SCSS	School of Computer Science and Statistics at TCD
SE	A gLite Storage Element

SQL	Structured Query Language
SSH	Secure Shell, a UNIX secure command line environment
StratusLab	An EU FP7 project aiming to integrate Grids and Clouds
SuGI	A Grid portal
TCD	Trinity College Dublin
t-Infrastructure	Training Infrastructure
Tomcat	An Apache servlet container
TransDeploy	A transactional deployment system
t-test	A statistical test of significance for differences between two groups
UI	A gLite User Interface
UML	User Mode Linux, a virtualisation technology
UNICORE	A grid middleware originating in Germany
UNIX	A multi-user operating system originating in Bell Laboratries
VARK	Visual, Aural, Read/write and Kinaesthetic learning style categorisation
Vivio	A bidirectional animation tool
VLE	Virtual Learning Environment
VM	Virtual Machine
VMWare	A virtualisation technology
VO	Virtual Organisation
VOCE	The central-European production Grid service
VOMS	Virtual Organisation Management System
VPN	Virtual Private Network
WebCom	An implementation of the condensed graphs model of computing
WebCom-G	Grid-enabled implementation of Webcom

WebCT	A learning management system
Weka	An open source statistical data mining application
Wiki	A web-based collaborative documentation tool
Windows	A proprietary operating system from Microsoft
WN	A gLite Worker Node
WS-Pgrade	A new web-services version of the P-Grade portal
X.509	A standard for certificates for a public key infrastructure
X86	The architecture of the Intel x8086-derived microprocessors
XeL	Executable e-Learning
Xen	A virtualisation technology
XMLDB	XML database format
XMLDB	Extensible Markup Language
Yaim	A fabric management system
youTube	A web tool for dissemination of video clips

Chapter 1

Introduction

1.1 Motivation

Grid technologies are maturing and entering a production phase, but outside of the High Energy Physics (HEP) community, there is still a dearth of knowledgeable users who can exploit the potential of Grids. To grow the Grid user community will require significant investment in training and education. This need has been recognised by the Grid community itself, and many recent projects, such as EGEE [1] and ICEAGE [2] have contained a strong education and training component.

Major policy and standards bodies such as the e-Infrastructures Reflection Group (e-IRG) [3] and the Open Grid Forum (OGF) Education and Training Community Group (ET-CG) [4] (of which the author was Secretary from its inception in 2006 until June 2008) have published documents in recent years calling for a sustained increase in funding and effort in Grid education and training.

These bodies have identified a skills shortage in research and industry which has the potential to cause existing infrastructures to be underutilised and to slow the emergence of knowledge-based economies [5], [6].

They argue that as distributed computing becomes more and more pervasive in our lives the need for trained professionals who are able to manage and troubleshoot these systems will grow rapidly [7].

eLearning seems ideally suited to Grid training and education because of the distributed nature of the Grid and its users, and their asynchronous training demand. eLearning allows learners to begin their training when and where they wish, while simultaneously reducing the overheads associated with face-to-face learning (organisation, catering, space, registration, resources, etc.).

Many Grid projects already provide online tutorials, most often taking the form of a static html user guide or a wiki. By only using these types of resources, however, projects and their students are lagging behind current eLearning practice. There have been many advances in the eLearning field in the last ten years and there are now a wealth of tools available for creating and delivering innovative online courses which follow modern pedagogic approaches such as constructivism, experiential learning, collaborative learning, etc. [8]

Furthermore, the emergence of adaptive courseware which can change details of the course, such as the exact content displayed, the difficulty level of the exercises, or even the entire pedagogic approach, mean that many of the disadvantages of eLearning as compared with face-to-face learning (for example, the one-size-fits-all approach and lack of feedback to the learner on their progress) are now being overcome. [9]

1.2 Objective and Goals

In this thesis I investigate the potential for using adaptive eLearning technologies to enhance Grid education.

The objective is to look at the questions of whether a learner's practical exercises running on the Grid can feed back into the personalisation of adaptive Grid eLearning courses, and whether this gives any benefit to the learner in terms of their learning experience.

In order to answer these questions, a secondary objective of this thesis is to attempt to prototype an appropriately integrated adaptive Grid eLearning system, and evaluate its usage by real learners.

Adaptive eLearning allows courses to be personalised for each individual learner, to provide a learning experience tailored to the particular user. This can help overcome some of the traditional problems of eLearning such as the one-size-fits-all approach and the lack of personalised feedback. Courses can be adapted based on a range of different characteristics of the learner, such as their prior knowledge, their preferred learning style or the characteristics of the client device with which they access the learning content.

Often prior knowledge and understanding of course concepts is tested in adaptive eLearning systems via multiple-choice tests. Alternatively, adaptive simulations can allow the learner to apply what they have learned in a simulated environment and the results can be fed back into the personalisation to update the prior-knowledge attributes of the learner's profile. This is a much better way of testing a learner's understanding of the material presented, but it is also considerably more difficult to achieve.

In the case of technology training there is a potential to go beyond simulation and to have the learner apply their knowledge in what is essentially a real environment.

This thesis looks at this idea of adaptive execution using a fair replica, that is to say, a replica that while perhaps not exact, is as similar as possible to and captures all of the salient features of the Grid-Ireland infrastructure [10] as an execution environment for learner's practical exercises. It investigates ways to use plugins and standard Grid services such as the Information System to feed into the eLearning personalisation and also ways to use the output of learners' jobs to generate further examples and exercises.

It explores whether the users' learning tasks (job submissions, etc) can actually run in the same environment in which the eLearning software runs, and if so, how well, and what issues arise, and whether as a consequence, the results can be made available to the eLearning software for use as inputs to the adaptivity engine and whether it can then modify subsequent instruction or exercises.

I have called this concept "Executable eLearning" (XeL). The central hypothesis is that XeL can enable practical exercises within the same environment as the adaptive engine, and that the synergy with and exploitation of XeL by the adaptive engine can significantly enhance Grid eLearning.

This thesis describes an architecture for an XeL-enabled adaptive Grid eLearning system consisting of

(1) a training infrastructure (henceforth referred to as "t-Infrastructure")

(2) an adaptive eLearning engine

(3) a web-based eLearning application which leverages the personalisation afforded by the adaptive engine and integrates with the t-Infrastructure

(4) a set of XeL modules to collect information about the results of learners' practical exercises and feed these into the adaptive engine

A prototype implementation of this architecture is described. This was implemented as an instrument to gather data to help answer the questions of this thesis, including those related to feasibility. In prototyping this instrument, the author attempted to use existing, open-source components where possible and to follow relevant best-practice in each area.

A number of controlled experiments were then conducted using the instrument in order to assess the central hypothesis.

1.3 Thesis Organisation

This thesis looks at the question of whether executable eLearning (XeL) can enhance the learning experience for Grid and distributed computing learners.

The thesis begins with a chapter describing a face-to-face course run by Grid-Ireland in 2006, the requirements identified from this experience, and how these requirements might be met by an eLearning approach.

This is followed by a review of the State of the Art in educational practice, eLearning and Grid and distributed computing education.

The next chapter describes a possible architecture of an XeL-enabled eLearning system as well as the prototype infrastructure (eLGrid) developed to evaluate the XeL approach.

Chapters 5 and 6 describe an evaluation of the XeL approach based on experiments conducted using the prototype tools.

Finally chapter 7 draws some conclusions and looks at directions of future work.

Chapter 2

An eLearning approach for Grid education and training

Adopting eLearning for Grid education has the potential to bring many benefits, and I discuss some of these benefits in this chapter. I also look at the rationale behind integrating the eLearning tools closely with the t-Infrastructure and choosing adaptive eLearning tools over the many non-adaptive tools which are available. I go on to describe the potential benefits which XeL might provide to learners in the Grid community.

The prototype system described in this thesis provides many of the benefits over traditional face-to-face learning. It also specifically addresses requirements which were identified as a result of a face-to-face course run in March 2006 by the Irish National Grid Infrastructure, Grid-Ireland. Section 2.1 describes this course, the requirements which arose from it and how the eLGrid system satisfies these requirements.

2.1 User Requirements

In March of 2006, the Irish National Grid Infrastructure, Grid-Ireland [11] ran a face-to-face training course entitled “Introduction to Developing Grid Applications” [12]. This was well-attended by Irish scientists interested in running their jobs on the

Grid. The course modules included a general introduction to the Grid and to the Grid-Ireland facilities, how to run jobs on the Grid, data handling on the Grid and running MPI jobs on the Grid. Hands-on sessions allowed learners to try out what they had learned.

At the end of the event participants were asked to fill in a feedback questionnaire. The questionnaire included evaluation questions along with spaces for learner comments.

In all, 16 feedback forms were returned by participants. A qualitative analysis of the comments was performed, with similar comments grouped together to try to identify themes which could be developed into a set of requirements for Grid education. The resultant requirements are discussed below.

2.1.1 Requirement 1: Ability to work at their own pace

The prior-experience and knowledge of participants in the face-to-face course varied. While some found the material easy, others found it difficult. Similarly, some found the sessions too long while others called for more time to be provided, particularly for the practicals.

- "Perhaps too much to take in in such a short space of time (for a beginner to the subject)"
- "Found the pace a bit slow"
- "Lectures felt very long"
- "Some parts a bit rushed, could have had longer practical on day 1"

Clearly, it seemed that participants would benefit from being able to work at their own pace.

Conclusion: eLearning is ideally suited to allowing learners to learn at their own pace. As outlined in section 2.2.1, it allows learners to access training when and where they want without the constraints of having to travel and follow the schedule of a face-to-face course. Instead, learners can fit their learning into their own schedule, covering as much or as little material as they choose at any given sitting.

2.1.2 Requirement 2: Certificates should be obtained in advance

A number of comments specifically asked for certificates (more properly, authentication and authorisation) to be issued in advance in order to avoid delays during the course.

- "Grid certificates should have been obtained before the class to maximise the time for the practicals"
- "Ideally all the work related to issuing certificates and VO membership should be done in advance"
- "If certificates could have been created in advance, and not during first half of the MPI Grid practical, it would have been better"

Grid authentication and authorisation is achieved using digital certificates issued by a Grid Certification Authority (CA). Each National Grid Infrastructure runs its own CA and in order to ensure the security of the infrastructure there are strict policies regulating the issuing of certificates to users [13]. The normal procedure for obtaining a Grid certificate involves visiting your local Grid Registration Authority (RA) and providing proof of identity such as a passport or staff or student identification card of your research institution.

This process does not lend itself well to courses, as learners must either apply for their Grid certificate in advance and bring it with them to the course, or the certificates must be issued on the day of the course once the local RA has verified the identity of each learner.

<u>Conclusion</u>*: These issues must be dealt with carefully in the case of eLearning; the problems may or may not be amplified depending on the approach taken. The infrastructure and instruments developed as part of this thesis must find a light-weight approach to issuing Grid certificates in order to ensure that learners can easily use the system. A t-Infrastructure with its own CA, which would issue short term certificates only valid for use on the t-Infrastructure would meet this requirement. This would*

allow certificates to be issued without following the rigorous policies of the National Grid Infrastructure. A consequence of the less rigorous security policies is that the t-Infrastructure would need to be designed for simple press-button re-installation in the event of inadvertent or malicious degradation. In addition, to further simplify the process for learners, it is likely to be helpful if the first course published in the prototype eLearning system were a course on "Obtaining a Grid Certificate" which took learners through the process of obtaining Grid credentials before progressing to the other courses, or indeed, before attending a face-to-face course.

2.1.3 Requirement 3: Practical environments and exercises

Eight of the 13 respondents who included comments in their response mentioned practicals. There was one general comment about the usefulness of being able to try out job submission while four participants mentioned practicals in response to the question "What did you most like about this event". Two respondents wanted more practicals, and another two mentioned problems in running the practicals.

The number of participants who referred to the practical sessions in one way or another suggests that the practical sessions are important to the success of any learning event.

At a face-to-face learning course, however, time constraints mean that it is not always possible to include as many practical exercises as tutors or learners might like.

Conclusion: *An eLearning course can present a large number of practical exercises which the learners are free to work on whenever they choose. They can also return to these as examples and reference materials when they have to implement their own, real-world solutions. An adaptive eLearning system can enhance this by identifying the most relevant practical exercises for each learner, based on their prior-knowledge and progress through the course, their particular learning goals, or other characteristics.*

2.1.4 Requirement 4: Personal support and feedback

The presence of tutors to answer questions and provide feedback was greatly appreciated by the participants.

- "Presenters and assistants were very patient and very helpful"
- "Ability to ask questions and depth of knowledge of the people hosting event"
- "This event give me the chance to familiar the most important command with the help of the expert" [sic]

Conclusion: *There are obvious disadvantages to eLearning when it comes to providing personal support and feedback to learners. One possible way to mitigate this problem is by using adaptive eLearning [14]. Adaptive eLearning systems provide personalised courses to learners, modifying content or navigation and providing feedback, based on characteristics of the individual learner, and on their progress through the course. Sections 2.2.3 to 2.2.4 look in more detail at the ways in which adaptive eLearning can address these issues, and whether other benefits can accrue from using adaptive and personalised learning.*

2.2 Benefits of a Grid eLearning System

2.2.1 Why eLearning?

The training investment needed to bring Grid usage to the wider scientific community is not insignificant and few organisations have the resources to keep up with demand. Traditional face-to-face training has significant overheads (travel of lecturers and students, arranging location, catering, registration, etc.). Furthermore, the workload is similar each time a course is run. eLearning, however, can allow a decrease in trainer-workload after the initial course creation. A well designed eLearning course which replaces the trainer interaction with technology allows the course to be given again and again without the overheads of a face-to-face course and with minimal effort on the part of the trainers (of course it may be necessary to have some interaction, regular updates to the course material, answering student queries, etc.).

In addition to the reduction in trainer workload, eLearning can help with the problem of asynchronous training demand. Training is required at different times for different users and it is often not feasible to run new face-to-face training courses for individuals or small groups as the need arises. Expecting users to wait until there are sufficient learners with the same need to run a face-to-face course is not satisfactory. eLearning allows new users to begin their training when they want it, not when it is convenient for the trainers, and to fit the training into their own schedule.

Even experienced Grid users will have need for ongoing training as new middleware versions are released, or to refresh their skills on a technology which they use infrequently. An eLearning system is ideal for these users who may not wish to attend a full training course. A well designed eLearning system could even alert users to changes and additions and recommend revisiting a previously completed course, supporting continuous professional development.

The arguments above can be made about any training activity, but in the Grid field there are additional reasons for considering eLearning. Grid education is particularly well suited to eLearning because of the distributed nature of the Grid and its users, providers and operators. When dealing with such a widely distributed user group it makes sense to design training which can be accessed from many and widely dispersed locations rather than to require users to travel to a central location.

More generally, eLearning is flexible and can be used both for self-paced learning and within the context of other training, such as face-to-face courses. Self-paced learning material delivered via an eLearning system can easily and effectively be incorporated into face-to-face courses. Indeed this is the direction which projects such as EGEE I, II and III [1] followed during the life of this thesis, and in which EGI [15] is now moving. Rather than having material presented by an instructor using power-point slides, introducing learners to the equivalent material via eLearning within a face-to-face course may have the added benefit of familiarising them with the eLearning system so that they may be more likely to use it to continue their training after the face-to-face course is over.

eLearning is not, of course, a panacea. There are some areas where traditional face-to-face learning outperforms eLearning, such as the interaction and personalised feedback from a tutor, peer interaction and support from fellow students, etc. These problems can be mitigated by careful design of the eLearning system, and incorporation of social interaction elements such as email, forums and chat.

2.2.2 Why Adaptive eLearning?

Many Grid projects already provide online tutorials, often taking the form of a static HTML user guide or a wiki. By only using these types of resources, however, projects and their students are lagging behind, and hence not taking advantage of, current eLearning practice. There have been many advances in the eLearning field in the last ten years and there are now a wealth of tools available for creating and delivering innovative online courses which follow modern pedagogic approaches such as constructivism, experiential learning, collaborative learning, etc. [8]

Furthermore, the emergence of adaptive courseware which can adapt to a learner's ability and progress by changing details of the course, such as the exact content displayed, the prerequisites covered, the level of difficulty of the exercises, or even the entire pedagogic approach, mean that many of the disadvantages of eLearning as compared with face-to-face learning (for example, the one-size-fits-all approach and lack of feedback to the learner on their progress) are now being overcome. [9]

The future of eLearning is likely to involve well designed courses tailored to the particular learner's requirements, providing feedback and varying the difficulty and content based on the user's ability and performance in online assessment. These adaptive tools are on the cutting edge of eLearning research today and should give a much more satisfactory learning experience.

When it comes to Grid education there is a large geographically distributed cohort of learners. Moreover, that cohort of learners is very heterogeneous, coming from diverse backgrounds with great variation in experience and prior-knowledge. It is not always feasible to provide face-to-face courses or have a trainer or tutor available to give individual support. Adaptive eLearning is ideally suited to these conditions.

2.2.3 Why integrate eLearning tools with t-Infrastructure?

Chapter 1 briefly outlined the concept of executable eLearning (XeL). This section and section 2.2.4 effectively ask "Why XeL?".

Where eLearning tools are to be used there are many compelling reasons to closely integrate them with the t-Infrastructure.

Self-paced practical exercises can easily be delivered by an eLearning system. But while the instructions for these exercises are delivered by the eLearning system, in most cases the exercises themselves are performed on the t-Infrastructure outside of the eLearning system. The first core concept of XeL is to perform the exercises from within the eLearning system, i.e. to integrate the eLearning system with the t-Infrastructure.

Several studies have shown that practice sessions distributed throughout a lesson lead to better retention of material [16] [17] when compared with training which places all of the practical exercises at the end of the lesson. In an eLearning course, however, distributing practical exercises throughout the lesson can lead to the learner chopping and changing between the eLearning environment and the training infrastructure where they must try out what they have learned. It is thus desirable to have the eLearning system integrated as closely as possible with the training environment. For example, if a web portal is being used for the practical exercises then it would helpful if this can be opened <u>within</u> a web-based eLearning page as a frame or via a new browser tab.

This allows instructions to be kept close to the exercise on-screen. While there is little in the research literature that looks directly at the benefits of this, studies have shown that retention and understanding are improved where text and images or other related objects are placed close to each other on-screen, the so-called "contiguity principle" [18].

If using adaptive courseware, then integrating the practical exercises into the rest of the course material would allow leveraging of the benefits of adaptation based on the learner's prior-knowledge, performance, or other characteristics which otherwise could only be applied to the didactic part of the course.

A more concrete benefit of this integration would be the ability of the eLearning tool to capture relevant information about learners' conduct of practical exercises, for example the inputs to the practicals, the behaviour of the learners when using the practical environment, the results of the practical exercises, etc. This information could then be used for assessment of the learner or adaptation of the course based on their results.

Integration would also allow leveraging of the single-sign-on capabilities of the Grid. If the eLearning tools can be modified to allow them to use the user's x.509 Distinguished Name (DN) to identify the user then the user need not log in separately to the eLearning tool. This will be even more useful with the new federated identity schemes such as EduGAIN [19] now being deployed which could automatically provide the whole of academia with privileges to access eLGrid.

2.2.4 Why adaptive executable eLearning

Chapter 5 describes in some detail the results of the experiments conducted with the eLGrid system before XeL is implemented. Learners were generally positive about the utility of the adaptive navigation in the form of a *"traffic-light"* annotated menu system. This suggests that the adaptivity (as implemented) benefits the learner.

This personalisation used the learners' multiple-choice test responses. However, such tests are not a good measure of learners real learning [20]. Practical exercises are a much better measure in courses that teach technology. If it is possible to link the practical environment and the eLearning tools, then this might make it possible to also adapt the learning material to the learner's conduct of practical exercises. This is the second core concept of XeL.

This thesis investigates whether any improvement can be found for learning achieved and satisfaction reported by learner groups who experienced this additional personalisation based on their practical exercises.

Furthermore, this thesis will also investigate the possibility of using the learner's performance on practical exercises as inputs to subsequent exercises, examples, etc. This would create the potential to enhance actual learning by providing more relevant examples.

If the learner's performance on the practical exercises can be made available to the eLearning software, the satisfactory completion of these exercises can be seen as prerequisites to completion of further courses, i.e. the concept of "prerequisites" will have been extended to practical exercises rather than just theoretical material. Any given concept could then be considered to be "understood" in several different ways, with a mixture of tests and practical exercises being used to determine the learner's prior knowledge, and their progress through the course material. This appears to be a closer approximation to instructor-led training than is normally possible with eLearning.

If it is so, then ideally course developers should be able to include prerequisite material, including practical aspects, into advanced courses on the basis that existing knowledge will be automatically taken into account, i.e. that adaptation will remove material, including exercises, that students have already learnt.

Section 2.2.1 suggests that learners may revisit the eLearning system in order to brush up or further explore topics of interest. Without XeL this relates more to the theoretical aspects of a course, although some t-Infrastructures allow learners to independently revisit practical exercises too. With XeL, the practical exercises would be able to be seamlessly revisited within the same eLearning environment, and adapted to the learners' current knowledge. Moreover, an integrated eLearning system might alert users about updates to the infrastructure and software of the training environment as well as course content modifications, allowing learners to revisit the practical components of courses and constantly update their skills for each new software release while easing the burden of updates for the course designers.

2.3 Future Potential of XeL

Modern eLearning tools are focussing on the support of constructive and experiential learning [8]. They are becoming more and more collaborative too [8]. An integrated practical environment with adaptive eLearning and collaborative tools might become the basis for a desktop problem-solving "personal XeL workbench" that includes learning and practical exercises adapted to existing knowledge, and where the t-Infrastructure allows speculative experiments in order to make progress towards a solution to a problem. The collaborative nature of the Grid enables natural extension towards a "personal/group XeL workbench", potentially bringing together groups of geographically distributed participants for distributed problem solving, where if the participants encounter aspects of a problem that are unfamiliar, then XeL is there to enlighten them.

While these predictions may be "visionary" wishful thinking, the perfectly reasonable theme is that XeL is not just a limited enhancement, but instead is a concept that creates the potential for extensive future development and impact.

2.4 Summary of Potential Attributes of XeL

The sections above discuss the various benefits of an XeL approach, and suggest several attributes of XeL. This section attempts to summarise these features:

1. Distributed pedagogically-sound self-paced adaptive Grid education via eLearning
2. Support for integration into face-to-face courses
3. Integration of t-Infrastructure for practical exercises <u>within</u> the eLearning environment
4. Capture of relevant information about practical exercises by the eLearning tool
5. Leveraging of single sign-on capabilities of the Grid
6. Adaptation to the information about learners' conduct of practical exercises
7. Ability to treat satisfactory completion of practical exercises as pre-requisites and for assessment purposes

8. Ability to revisit adaptive eLearning including integrated practical exercises

9. Possible extension to become an integrated problem-solving "personal/group XeL work-bench"

Of these, (3), (4) and (6) are the core concepts of XeL.

Chapter 3

Current State of the Relevant Art

3.1 Introduction

It is impossible to begin to discuss eLearning without proper reference to the fundamental learning paradigms and theories which aim to explain how humans learn. Thus this literature review begins with a discussion of some of the more mainstream learning theories of relevance to this thesis.

Similarly, there has been an explosion in eLearning and adaptive eLearning research in the last three decades and so there follows a discussion of a number of the relevant approaches and theories.

The state of the art in eLearning as it relates to Grids and this thesis also encompasses t-Infrastructures, Grid education and Grid eLearning. This chapter then looks at some of the most relevant work in each of these fields.

The discussion section reviews the related work in light of the thesis goals and attempts to set some design guidelines for the development of an XeL architecture as well as the implementation of the prototype instrument.

3.2 Learning Theory

There are many learning theories which aim to explain the processes at work in acquiring or modifying knowledge and skills. For a comprehensive review, see for example [21].

Behaviourist approaches, firmly situated in the empiricist epistemology tradition, argue that the mind, which cannot be observed, is therefore not a suitable object of study. They view learning as essentially a process of modification and acquisition of behaviours, where learning can thus be objectively measured through observation of behavioural change. Similarly, as learning occurs via external stimuli working upon the learner to reinforce some behaviours and discourage others, the role of the teacher becomes one of presenting the learner with appropriate learning resources and tasks (the stimuli) which bring about changes in behaviour (the response) [22].

Cognitive theorists view the behavioural approach as simplistic, dealing only with observed behaviour while ignoring the cognitive processes which lie behind this. They argue that learning cannot be understood without reference to these cognitive processes [23]. Within the cognitive approach there are a wealth of theories investigating different aspects of the cognitive processes associated with learning.

Cognitive theorists believe that knowledge and understanding are actively constructed by the learner, not passively received from the teacher (as is the case in behaviourist theories of learning). Knowledge is represented in the brain in the form of mental models or schemata, and these can be modelled and studied [21].

Cognitivists also take account of the existing preconceptions and ideas of learners which may influence their receptiveness to new knowledge either positively, by giving them an existing model into which they can incorporate new knowledge, or in the case of misconceptions, negatively, potentially militating against the acquisition of new knowledge. Learners construct knowledge through interaction with the physical world.

Constructivism views learning as the process of actively constructing schemata. It builds upon cognitivism, and also stresses the importance of understanding the cognitive processes behind learning. It includes similar ideas about mental models, schemata, preconceptions and learning styles.

Although some critics argue that the term is ill-defined and too broadly applied to a wide range of different approaches [24], constructivism is widely considered the dominant theoretical framework for the current understanding of education, in particular, science education [24].

Each of these approaches has some relevance to this thesis. In particular, many theories from the cognitive and constructivist approaches can be fruitfully applied to Grid eLearning.

3.2.1 Specific theories of learning

A number of theories are of particular relevance to the work of this thesis. These are largely within the cognitive and constructivist paradigms, but some also stem from the behaviourist approach. This section describes some of these theories.

Active learning

Active learning theory follows the constructivist approach and emphasises the learners' role in the learning process, arguing that learning is more effective when the learner takes an active role in the learning process. It stresses the importance of activities that promote higher-order thinking tasks, such as reflecting on learning, analysing and evaluating. Group discussions, writing assignments, case-studies, debates, role-playing and simulations are common methods employed in active learning [25].

Recommendations for eLearning systems based on active learning include providing many opportunities for practice, and engagement with real-world case problems and examples to encourage the learner to be cognitively active. However, some research has shown that being too active, in particular, being behaviourally active, can in fact reduce successful schema building and thus learning [26].

Another way to encourage learners to be appropriately cognitively active is to provide opportunities for them to review and reflect on their own learning. For example by showing their progress through the course material and to provide feedback on the process as well as the products of learner problem-solving, or by explicitly providing opportunities for learners to review and document their problem-solving techniques and to compare these with those of other learners and of experts [18].

Active learning techniques unique to eLearning include practice with automated tailored feedback and the use of simulation to accelerate expertise [18].

Problem-based learning

Problem-based learning is another active learning technique which involves introducing a problem at the beginning of a lesson and encouraging learners to choose what to learn in order to solve the problem [27]. It is argued that this helps learners to build schemata as they are not acquiring abstract knowledge without an understanding of its applications.

An eLearning course can easily support problem-based learning, providing opportunities for learners to perform real-world problem-solving tasks. Learning materials can be structured such that the problem is presented at the beginning, and the learner can then browse the learning material, perhaps freely, or with guidance, in order to find the solution. In an adaptive eLearning system, the content can even be adaptively presented in a problem-based or traditional didactic manner depending on the attributes of the learner, for an example see [28].

Worked-example Effect

Worked-examples also fall under the 'active-learning' banner, but unlike problem-based or discovery learning, the learner is not actively solving a problem. Instead they are cognitively active in understanding the problem being solved. This has been shown to be more beneficial in some cases than active problem solving [29].

Worked examples allow expert problem-solving actions and thinking processes to be explicitly presented and explained to the learner, and this can be beneficial particularly to novice learners. In an eLearning course it may be advisable to replace some practice problems with worked examples for novice learners, and intersperse worked examples among practice problems [18], and here again, there is the potential for an adaptive eLearning course to present worked examples or problems to the learner based on their prior knowledge and expertise.

The importance of Practice

In order to master any skill it is necessary to practice. While this may, in many ways sound like simple common-sense, a recent opinion-piece [30] by Tim O'Reilly, of O'Reilly Books, noted that while practice is essential when learning technology subjects such as programming, most learners will try to do the minimum number of exercises possible.

In an eLearning environment where learner drop-out is usually high and it is generally hard to ascertain the level of learner engagement with the material, it becomes necessary to carefully design the system, materials and exercises in order to encourage the maximum participation in the practical exercises.

In order to ensure that the practical exercises are interesting to learners and thus reduce drop-out rates, exercises should require learners to apply knowledge and skills to real-life situations [18].

Furthermore, research has shown that distributed practice exercises are significantly more effective and contribute more to learning than massed practice, for example, where all the practical exercises are grouped together at the end of a course or after an initial lesson [31].

The implications of such studies for eLearning are clear. It is desirable to have several practice exercises per topic, distributed throughout the lessons, and they should be based on real-world problems which are relevant and interesting to the learner.

Cognitive Load theory

Cognitive load theory is a theory of learning from the cognitive approach which deals with attention and the transfer of skills and knowledge from short-term to long-term memory.

The theory states that learning occurs when a change is effected to long-term memory, or to the mental schema held in long-term memory. Rehearsal can allow the recall and use of these schemas to become automated, that is, the knowledge and skills can be used without the learner having to pay attention.

In order to alter long-term memory, however, novel information must first be processed in short-term or working memory, which is severely constrained in terms of its capacity.

Cognitive load is the load placed on working memory during learning. It is divided into three types

- *Intrinsic* load due to the complexity of the subject matter being learned
- *Extraneous* load due to instructional design which fails to take account of the limits of working memory
- *Germane* load due to the active learning effort of the learner.

This theory has important implications for multimedia design and eLearning in terms of the organisation and presentation of material, how much novel material should be presented and the mode in which to present novel material [32].

For example, the *split-attention* or *contiguity principle* argues that where a learner must integrate information from different sources, placing them close together on screen reduces cognitive load. Practical recommendations which can be drawn from this include placing text close to associated images, and instructions for practical exercises close to the interface where the practical steps are carried out [18].

The *modality effect* is based on research which shows that each input "channel", corresponding to the different human senses, has a separate short-term memory store. This theory argues that one should present different sources of information in different modes (visual, audio, dual-channel) [33]. For example, one might use audio narration rather than text to explain on-screen graphics or animations.

The *redundancy effect* states that redundant learning material decreases learning. Unnecessary repetition, or the presentation of the same material in different ways, can increase cognitive load. Furthermore, the *expertise-reversal effect* shows that the information necessary for a novice may well be redundant to an expert. Thus some of the techniques employed to take into account the contiguity principle or modality effect may actually end up increasing cognitive load for an expert learner [18].

While an experienced teacher or tutor can usually tailor the amount of material provided to the level of the learner, in an eLearning system this requires personalisation and adaptive eLearning technologies [34].

Worked examples can decrease cognitive load, but some active learning and problem-based learning techniques can actually increase extraneous cognitive load [29]. Problem-based learning and discovery learning require the learner to use some of their attention and short-term memory to identify what they should be learning, while with worked examples there is no search of the problem space and all attention can be given to the formation of new schema. A mix of problem solving and worked examples can help to balance these effects.

3.3 eLearning

In the broadest sense of the term, eLearning or *electronic learning*, is any electronic or technology assisted learning or teaching. The term, and indeed the entire field of eLearning, has developed out of research in the areas of computer-based training (CBT), distance learning, military simulations, and other fields, each with their own focus and approach, and thus it is difficult to pin down a concise definition of the term [35]. The modern usage in most sectors, however, emphasises internet-based learning experiences [35] and this is the sense in which the term is used in this thesis.

While earlier systems often focussed solely on content delivery and were little more than computerised books, modern eLearning systems include a diverse range of multimedia, interactivity, social learning and support for learner-centred pedagogy.

Virtual Learning Environments (VLEs) also known as Learning Management Systems (LMSs) such as Blackboard [36] or its open-source competitors Moodle [37] and Sakai [38] are now commonplace in tertiary education [39]. These facilitate content delivery and usually also provide some mechanism for student-student interaction, and student-tutor interaction, as well as support for assessment in the form of assignment submission and automated multiple-choice assessment tests [40]. These systems are also widely used in industry to support training of new staff, continuous professional development and compliance certification [41] [42].

While several research prototypes have attempted to add personalisation to VLEs (for example [43], [44], [45]), VLEs generally do not provide much built-in support for personalisation or adaptivity beyond basic localisation options.

This section identifies recent eLearning research areas which are of particular relevance to this thesis, and reviews similar work to the adaptive executable eLearning framework proposed herein.

3.4 Adaptive eLearning

Adaptive eLearning allows courses to be personalised for each individual learner, with this personalisation potentially taking many forms. Common personalisation techniques involve modifying the navigational path that the user takes through the course (adaptive navigation) or modifying the learning content to the user (adaptive presentation).

Similarly, the adaptation can be based on a range of different characteristics of the learner, such as their prior-knowledge, their preferred learning style or the characteristics of the client device with which they access the learning content.

3.4.1 Methods of Adaptivity

The way in which the course is personalised is the method of adaptivity. This generally involves modifying details of the course content, structure or the learning environment.

Adaptive Presentation

Adaptive presentation involves presenting different material, or the same material in different ways, to different learners. This might include:

- Presenting localised content which takes account of the learner's language or other region-specific details, or localising the examples given to make them more relevant to learners from a particular region, for example [46] (see section 3.4.2)

- Varying the difficulty of the material presented based on the learner's ability, prior knowledge or expertise [47] (see section 3.4.2)

- Changing the presentation media to suit learners' style, for example, based on the learner's VARK (Visual, Aural, Read/write and Kinaesthetic) style [48] (see section 3.4.2)

- Providing low- and high-bandwidth content based on the learners' connection speed [49] (see section 3.4.2)

Adaptive Navigation

Another option is to adapt the navigation path through the course. This can involve any of the following:

- Link Annotation: The addition of extra information which provides the learner with some information about the status of the content linked (for example, its relevance to their learner profile, or whether they are ready to view it based on their prior-knowledge) [14].

- Link-hiding, disabling and removal: Links not relevant to a learner might be disabled, so that the link anchor appears as ordinary text, or the link may be hidden or removed from the navigation menu [50].

- Re-ordering concepts: The ideal order in which concepts should be tackled may differ for different types of learner. One might choose, for example, based on the Honey and Mumford learning style categories, to present theory first to a theorist learner or the problem first to an activist learner [51]. For an example of a system which does this see [28].

Various types of navigation control can be supported in an adaptive eLearning course ranging from *learner control*, a discovery-based learning approach, through to *program control* where the learner is given no option to influence the order in which subjects are tackled [18]. In between these extremes are various forms of advisement.

Learner control can be effective where the content is relatively uncomplicated and there are few interrelations between topics, or where the audience has significant prior-knowledge of the subject area. Program control is suitable mainly for novice users. Advisement is a more flexible approach which can be used where the learners have a mix of knowledge and skills or where saving time of those users who can demonstrate prior competence is a priority [18].

3.4.2 What to adapt to

Adaptation to the learner's style

Different learners have different learning strategies and styles, and benefit to different extents from particular structuring or presentation of materials. There is a lot of research in this area and several different theories attempt to classify learners into different groups based on their learning style.

For example the VARK (Visual, Aural, Read/write and Kinaesthetic) [48] (see also [52]) categorisation developed by Neil Fleming of Lincoln University, New Zealand, groups learners by the mode of delivery they prefer and argues that similar content should be presented in different ways for different types of learner.

Kolb [53], and Honey and Mumford [51], instead focus more on the underlying cognitive style of the learner. Kolb proposes two integral processes in learning: *prehension* of concepts and *transformation*. Each of these processes can take place in two different ways, prehension via *apprehension* (grasping through immediate concrete experience), or via *comprehension* (indirect grasping through symbolic representations of experience), and transformation via *intension* (reflection), or via *extension* (action).

These can be combined to give four styles of learning:

- *Divergent*: apprehension transformed through intension
- *Accommodative*: apprehension transformed through extension
- *Assimilative*: comprehension transformed through intension

- *Convergent*: comprehension transformed through extension

Honey and Mumford's theory of learning styles categorises learners into *Activist*, *Reflector*, *Theorist* and *Pragmatist* groups. While these are similar to Kolb's classifications the real strength of the Honey and Mumford approach is the manual published by the authors, which serves as a toolkit for educators, enabling them to easily implement learning-styles testing and giving a *cookbook* of ways in which the theory may be applied to eLearning activities. This makes it an attractive option for educators wishing to implement personalisation based on learning style and it has been used in many adaptive eLearning systems [28], [54].

Other approaches include Gardner's theory of multiple intelligences [55], Pask's Learning strategies [56] and Hudson's convergent and divergent thinking [57].

Learning styles are a popular element for adaptation in personalised eLearning courses and many systems or prototypes exist which implement personalisation based on different learning style theories. For a review of the literature in personalisation based on learning style see [58] or [59].

Such systems are generally based on the notion that a learner will learn best when content is presented to them in a way which matches their personal learning style. This might, for example, involve changing the type of content from plain text to graphics, audio or multi-media if using VARK, or, as another example, changing the content itself to be more theoretical for an *Assimilative* learner under Kolb's classification.

The course navigation can also be adapted based on learner style. In problem-based learning, expecting a learner to randomly explore a problem space can be inefficient [26], and this is likely to be even more pronounced where the learner's style is not suited to this approach. Thus adaptive navigation techniques could be used to help leverage the benefits of active learning and problem-based learning while not imposing excessive extraneous cognitive load. The degree to which the navigation is structured for the learner could vary based on the learner's style, for example, presenting a fully problem-based learning or discovery learning approach for learners who are *Activists* under the Honey and Mumford classification, while providing a more structured navigational path for *Theorist* learners.

Some systems also have a secondary aim of educating learners about their preferred learning style or training them to use learning styles other than that which they naturally prefer, thus teaching the learner to learn and improving their cognitive flexibility.

Existing systems usually determine the learner's preferred style via a pre-course questionnaire.

Adaptation to the learner's Devices

Adaptive eLearning systems often personalise course delivery based on characteristics of the learner's environment, such as the device with which they are accessing the course, or the bandwidth of their connection.

For example, a system might use adaptive presentation to display a low-resolution version of course materials to a learner using a slow or high-cost connection [60], [61], [49]. This could involve lower resolution images, text-based material instead of graphics, the exclusion of certain materials such as videos, etc.

Similarly, certain types of learning content might be more suitable for users connecting on devices with small screens such as mobile phones, for example a graphic might be simplified to ensure that it is visible on a small screen, or a landscape or portrait version might be used for different screen sizes if switching between the two orientations is not possible on the device [60]. Or, for example, the eLearning environment can be made more convenient for mobile users by integrating it within smart phones (for example, Intuition Publishing Ltd. [62] has a dominant role in provision of eLearning integrated with Blackberry devices).

Adaptation to the learner's Location

It may be necessary to adaptively deliver different content to users in different regions. This may be for cultural reasons [63] or for licensing purposes (e.g. where an application is licensed only for users in a particular location).

Adaptation to the Learners' Prior Knowledge and Expertise

Any concept within a course can be considered to be understood or not, and the navigation or content can be modified based on which course concepts are understood by each learner.

This is often assessed at the start of the course, for example via a pre-course learner questionnaire or test. The learner's profile can also be updated as they progress through the course, with a concept considered understood when the learner has viewed the relevant material, or after completion of multiple-choice test questions relating to that concept.

Automatic assessment of such multiple-choice tests is easy to achieve, hence non-adaptive test functionality is already built into many Virtual Learning Environments (VLEs). Similarly, there are several adaptive eLearning tools which modify aspects of a course based on the results of such tests, such as the Adaptive Personalised eLearning Service (APeLS) developed by the Knowledge and Data Engineering Group in Trinity College [64].

Multiple-choice tests, however, are not a very reliable way of measuring a learner's knowledge [20]. Adaptive simulations, described in section 3.4.3, can allow the learner to apply what they have learned in a simulated environment and the results can be fed back into the personalisation. This is a much better way of testing a learner's understanding of the material presented, but it is also considerably more difficult to achieve.

In the case of technology training there is a potential to go beyond simulation and to have the learner apply their knowledge in what is essentially a real environment which feeds back to the eLearning personalisation engine, and this is the approach taken in this thesis with XeL.

3.4.3 Simulations

Simulation in eLearing involves using technology to model an environment in which learners learn or practice new skills. Simulation is primarily concerned with learning skills rather than knowledge, that is *learning to do* rather than *learning to know* which has been the focus of the majority of traditional eLearning tools in the past [65].

Simulations can be used in many domains, from soft-skills training through computer skills to military and flight simulations or even medical diagnosis and surgical procedures. The Futurelab project has published a review highlighting many prominent examples of educational simulations [66].

There are many tools which enable the creation of software training simulations, for example using animations of sequences of mouse-clicks and keystrokes required to complete various tasks. This is often achieved by creating a flash recording of the tutor performing the same tasks using the real software. To this can be added some interactivity where the learner must select what to do next, or questions can be embedded in the simulation. Popular simulation systems supporting this type of screen-capture and replay include Adobe Captivate [67], Assima [68] and Xstream Rapidbuilder [69].

Other tools provide for different types of simulations, for example educational immersive games or softskills training [70], [71]. An interesting example from the author's host research group is Vivio [72], a bidirectional animation tool that allows simulations to move both forward and backward in time.

Integrating simulations and adaptive eLearning is difficult, however, and few systems attempt it. Those which do often embed the adaptivity directly into the simulation making it difficult to generalise or reuse the content. One ground-breaking tool for creating adaptive simulations is ACTSim [73] but this is focussed on soft skills training.

3.4.4 Assessment

Assessment is the process of evaluating the learner's progress and the extent of learning achieved. *Summative assessment* aims to establish a measure of the student's learning at the end of a course, usually in order to award a final mark or grade, while *formative assessment* is focussed more on establishing the student's progress during a course in order to provide feedback or take remedial intervention if learning progress is not in line with expectations [74].

Assessment is often supported by VLEs, which typically provide multiple-choice test functionality to test the knowledge of the learner and can sometimes provide limited feedback to the learner on their results. Multiple-choice questions are quite limited in their ability to test *deep* learning or understanding gained from a course. They also fail to assess whether the learner has been able to transfer the skills and knowledge gained to real world problems. For this type of assessment it is usually necessary to involve a tutor.

Short answer marking engines [75], which analyse free-text answers for keywords, can provide more realistic and flexible assessment, as can comparing answers or outputs of practicals with work of experts.

Computer-assisted marking can reduce the time taken to mark assignments, for example allowing the computer to mark certain portions of an assignment, or automatically clustering similar answers to which the tutor can easily assign a single grade. Assess by Computer (ABC) [76] used by the University of Manchester is an example of a computer-assisted marking application.

Simulations can enable closed-loop adaptive education which involves the use of automated adaptive assessment via practical exercises carried out in the simulated environment [77]. One of the proposed benefits of XeL and the integration of the t-Infrastructure with eLearning tools is that information about the learners conduct of practical exercises should become available to the eLearning system, and this may allow closed-loop adaptive Grid education.

3.5 Grid Education

The Grid field, and more broadly the e-Science and e-Infrastructure fields, have many initiatives related to education and training. This section reviews some of the most important activities of recent years.

The most celebrated educational events in the field are the International Summer/Winter Schools of Grid Computing (ISSGC/IWSGC) series [78]. The summer schools were run from 2003 until 2009, catering for approximately 60 - 80 students each year, most from Europe but substantial numbers also coming from the US, Asia and other regions. An online winter school was introduced in 2007 [79]. These events covered multiple Grid and distributed computing technologies from a user point-of-view and gave a grounding in the theory as well as practice of Grid Computing.

Another important Grid summer school in Europe is the GridKa Summer School [80], held each year at the Karlsruhe Institute of Technology (KIT) in Germany. The school holds parallel practical sessions in order to cater for learners from different backgrounds, with separate tracks for the end users of Grids (application scientists) and the system administrators and site managers who run the Grid infrastructures.

Open Science Grid [81] is the American Grid initiative equivalent to EGEE/EGI in Europe. OSG runs regular Grid Schools throughout the US each year. The focus is not as broad as the ISSGC event with only OSG technologies covered. In 2010 OSG organised its first Summer School, along the lines of the ISSGC, but retaining the narrower focus of the OSG Grid Schools.

The Distributed Computing Infrastructures projects (DCI), funded by the European Commission's Framework 7 programme [82], plan to hold the European Distributed Computing Infrastructures Summer School (EUDCISS) on Grid, Cloud and Desktop Computing Services [83] in July 2011 which will also take inspiration from the ISSGC events, but will focus specifically on those technologies developed by the DCI projects.

Several Grid user communities also run their own summer schools focused on using Grid technologies within their discipline. For example the BioMed Summer School [84] which was first held in 2007, or the CERN School of Computing [85] run since 2004.

While formal courses within Universities which include a Grid computing component are still relatively rare, there are a number of examples such as the *MSc in Grid Computing* at the University of Amsterdam [86], the *MSc in Distributed Scientific Computing* (previously *MSc in e-Science*) at the University of Edinburgh [87], or the *MSc in Distributed Computing Systems Engineering* at Brunel University [88]. The ICEAGE project [2] has a list of e-Science and Grid MSc courses on its webpage and the e-IRG has also created a wiki page to collate further examples of courses and other Grid training and education initiatives [89].

Here at Trinity College Dublin the course CS4021 (formerly 4BA9) *Advanced Computer Architecture* [90] has been teaching Grid computing topics to Computer Science undergraduates for some years now.

Many think-tanks, work-groups and other fora in which to discuss the issues of Grid and Distributed Computing education exist. These fora have made several important contributions in the area of policy and standards by publishing recommendations, reports and white papers to inform the work of Grid and distributed computing education professionals world-wide. Some examples are given below.

The Open Grid Forum (OGF) [91] for several years had an Education and Training Working Group (of which the author was secretary from its inception in 2006 until June 2008) which brought together experts in the field of Grid education. This group worked to catalogue initiatives in Grid education such as MSc courses, Grid training infrastructures, etc. They also published two OGF recommendations documents:

- Policy for Supporting Grid and e-Science Education and Training [6]
- Towards Professional Grid Certification [92]

The ICEAGE Forum was an initiative of the ICEAGE project [2] which brought together stakeholders in the field of Grid Education from beyond the project in order to provide input to the ICEAGE project and also to give the ICEAGE project's outputs an audience of high-level figures in the Grid Education field who could ensure that policy recommendations from the project were brought to the attention of the appropriate people and organisations within their own countries.

The ICEAGE project published several documents based on the discussions and recommendations of the ICEAGE Forum and various workshops organised by the project, such as the *Curricula for undergraduate and masters level courses in e-Science: Report from the ICEAGE Curricula Development Workshop* [93] which sets out a basic e-Science curriculum for graduate level courses covering core competencies such as critical thinking and digital systems thinking along with a solid base of statistics, numerical models and issues relating to data curation.

Other important contributions by the ICEAGE project are the ICEAGE Digital Library [94] and the Grid People Registry [95].

The e-Infrastructures Reflection Group (e-IRG) [96] is an EU body tasked with setting policy and direction for e-Infrastructures in Europe. The group's Education and Training Task Force (ETTF) produced a report in 2008 [7] and provides input to the e-IRG white papers [5]. These documents propose strategic actions to promote Grid and e-Science education in Europe. The ETTF report calls for a more sustained e-Science and e-Infrastructures education programme, covering many of the core competencies identified by the ICEAGE Curricula document, to be embedded into the normal academic programmes of member states. This would be complemented by training in specific technologies such as Grid computing.

These documents call for a more systematic approach to e-Science education including Grid, distributed systems, statistical thinking, and other topics. Of particular interest to this thesis are the repeated calls for provision of training infrastructures [5] and [6].

The e-IRG ETTF report also makes specific mention of the need for multiple modes of delivery for distributed computing education, which it regards as necessary: *"Different target audiences would require the presentation of different principles, concepts, and examples, so that the mode of delivery and curriculum are geared towards that audience. Flexible refresher courses could update students on new technologies and summer schools could appeal to academics who would not have time to commit to a Masters course in Grid computing or e-Science."* These needs could be addressed using adaptive eLearning, which would allow courses to be personalised for the various target audiences, and allow online refresher courses for Grid professionals and academics.

The European Grid Infrastructure (EGI) [15], the umbrella entity for the European National Grid Infrastructures (NGIs), has organised a Training Task Force, of which the author is a member, that aims to coordinate the training efforts of the NGIs, DCIs and other related Grid projects and encourage sharing and re-use of materials, infrastructures and expertise. The task force will reuse a number of services produced by the ICEAGE project, such as the Digital Library and the Grid People Registry.

The European Middleware Initiative [97] is a major supplier of software for the EGI and NGIs. The project is a collaboration of four major Grid middleware providers in Europe: ARC [98], dCache [99], gLite [100] and UNICORE [101]. The Technical Knowledge Transfer activity within EMI, which the author leads, is responsible for training members of the project and end-users on EMI technologies. It has organised a number of training events and is attempting to create a body of online training materials, the first steps of which have involved delivery of live training events in an online mode via the Adobe Connect video conferencing system [102]. Training sessions are recorded and made available for subsequent download via the EMI website. EMI is also investigating the use of tools such as the ICEAGE Digital Library.

The StratusLab project [103], in which the author is also a member, aims to integrate Grids and Clouds. The project has not yet created a large body of training material or held many training events, but has produced training videos [104] covering the use of their tools. These have been made available via *youTube* and could easily be embedded into any VLE or adaptive e-Learning system for reuse within courses covering cloud technologies.

More broadly, almost all Grid projects participate in some dissemination and training activities, and there are a number of projects which are mainly dissemination avenues, but which also serve an educational purpose in informing and educating about new developments in fields such as Grid Computing, e-Science and e-Infrastructures, High-Performance Computing and Cloud Computing. For example the e-Science Talk project (formerly GridTalk) [105] coordinates the dissemination outputs of EGI and other European e-Infrastructure projects. e-Science Talk publishes *International Science Grid this Week* (iSGTW) [106], an online magazine covering topics in Grid and scientific computing, and also maintains a number of other dissemination outlets such as GridGuide [107], GridCafe [108] and GridCast [109]. It also publishes GridBriefings [110], which is a periodic topic-based newsletter which summarises key developments in the field. In July 2010 GridBriefings published a briefing entitled *Putting the 'e' in education: eLearning and grid computing* [111] which discussed some of the benefits of and challenges for Grid eLearning, as well as providing a number of case studies and examples.

3.6 Grid eLearning

The International Winter school of Grid Computing mentioned in section 3.5 [79] is an example of a Grid education initiative delivered entirely online. The course is conducted using the Adobe Connect [102] web conferencing system and is run over a period of five weeks covering four technologies (Condor, OGSA-DAI, Globus and gLite). A number of live online events are held throughout the five weeks with asynchronous work in between. The live events include keynote talks by experts in the Grid field as well as tutorial sessions. Lecture notes and practical exercises are provided on a webpage and students work on these in their own time over the course of the school with chances to ask questions during the live tutorials, or via email.

Other projects have followed the example of the Winter School in providing "live" online events via video conferencing tools, for example the European Middleware Initiative [97] recently held a three day *in-reach* event (a training event aimed at members of the EMI project) [112] with live presentations and question and answer sessions using the Adobe Connect software.

Many projects provide wikis or web pages containing documentation as well as Powerpoint presentations from previous face-to-face courses, videos of talks and lectures and instructions for practical exercises. The GILDA training wiki [113] and the OSG Education, Outreach and Training wiki [114] are examples of such sites. While such systems are valuable resources for the Grid community, they make little attempt to present pedagogically-based instruction to the learner and are more akin to repositories of documentation, reference and training materials than true eLearning. The primary benefit of such systems may be in their value to expert learners who already have a high degree of familiarity with the technologies and underlying principles, but they are likely to be of less use to novice learners who require more pedagogic support and direction in their learning.

There are a small number of projects providing educational access to instrumentation Grids to allow learners to interactively control scientific instruments and perform their own virtual experiments, such as BugScope [115] or [116].

The ICEAGE Digital Library, mentioned in section 3.5, is a digital repository containing a huge number of educational resources including presentations, audio recordings of lectures, books, podcasts, tutorials, videos, exercises and assessments. These materials are freely available under the Creative Commons licence [117]. Educators can reuse and modify them for their own courses and for individual self-paced learners. The SuGI portal [118] is another repository of training materials.

Access Grid [119] is a virtual collaboration tool providing video conferencing facilities along with data sharing and access to Grid applications. It has been used in many educational contexts to allow collaborative project work [120], [121]. Grid technologies other than Access Grid can also be incorporated into collaborative eLearning tools, for examples see [122], [123] [124].

The European Learning Grid Infrastructure (ELeGI) [125] project, which ran from February 2004 to January 2008, aimed to use Semantic Grid technologies to facilitate a move from simple information transfer in eLearning systems towards a more learner-centred, constructive and experiential learning approach. The project's goal was to produce a semantic Grid for human learning: the Learning Grid. Learning would be provided in the form of Semantic Web or Grid services, and different combinations of these services would allow personalised learning to be delivered to learners.

3.7 Practical Environments for Grid

As discussed in section 3.2.1, practice is an essential part of learning any technological subject. In order to provide opportunities for practice of skills it is necessary to provide a practical environment where learners can try out what they learn.

This section looks at the state of the art in Grid learning practical environments, as well as interesting examples of other eLearning practical environments.

3.7.1 Grid t-Infrastructures

The provision of Training Infrastructures, or t-Infrastructures, is a core part of any Grid education and training plan. In order to effectively teach Grid technologies, practical experience is vital. Learners must be given a practical environment where they can try out what they learn and complete the practical exercises which form part of their courses.

An excellent and thorough overview of the issues involved in t-Infrastructure provision can be found in [126].

GILDA

The most widely used Grid t-Infrastructure is the *Grid INFN Virtual Laboratory for Dissemination Activities* (GILDA) [127] initiative of INFN [128] which began as a t-Infrastructure for the EGEE Grid [1] but has since gone on to address a range of different Grid middlewares. It is now used as a training infrastructure for EGI. Organisations participating in GILDA provide some portion of their Grid resources to act as dedicated training resources. GILDA issues short-lived certificates valid only on these resources to learners wishing to use the Grid. GILDA resources can be used at any time by individuals, or a subset of resources can be booked for use in courses.

The certificate issuing process is simple, with generic certificates being issued for a period of two weeks. Unlike when a real Grid certificate is issued, complicated security checks are not required (because of the limited scope of the credentials). This makes it easy for anyone to get a certificate and try out the Grid.

Almost 500 training and dissemination events have been supported by GILDA over the period 2004 to 2011 and over 14,000 personal certificates issued to learners [129].

GILDA has gradually moved towards a model of using virtual machines and now provides VM images for worker nodes, user interface machines, etc., so that sites can more easily join the GILDA testbed. These VMs can also be used in Grid administration courses where learners are required to set up their own GILDA sites.

Some generic Grid portals are able to harness GILDA:

- The Genius Portal [130] is a Graphical user interface to the Grid allowing for job submission and other actions. Genius is installed on the GILDA resources and so can be used by learners using their GILDA certificates to run jobs without having to be familiar with the UNIX command line interface.

- The P-Grade Portal [131] can be configured to use the underlying GILDA resources. P-Grade, however, is also an advanced workflow portal allowing users to construct complicated workflows composed of multiple services and to connect together their inputs and outputs. Any learner with a GILDA certificate can apply for a P-Grade account and use the system. P-Grade also supports accessing non-GILDA resources if the user has the appropriate certificate. Recently it has been largely re-written for web services; this new release is called WS-PGrade [132].

GILDA has been used by a number of projects for their training and education goals. It was the main t-Infrastructure used by EGEE-II and EGEE-III [1] and was also adopted by the ICEAGE project [2]. It is now used by the EGI-InSPIRE project [133] along with most of the European National Grid Initiatives (NGIs). GILDA is already used in tertiary education, particularly in Italy, where is has been used by a number of PhD students for their Research.

It is also used in many other countries through the EELA [134], EU-IndiaGRID [135], EUMEDGRID [136] and EUChinaGRID [137] projects.

Despite the success of GILDA there have been several other t-Infrastructures used by various national Grids over the past years. For example, the Central European Production Grid Service, VOCE [138], has used a separate training Virtual Organisation (VO) on their production infrastructure, while the Greek National Greek Research & Technology Network, GRNET [139], has used high-priority queues on their production infrastructure in order to ensure that learners receive timely feedback on their training jobs.

Open Science Grid

The OSG maintains a small number of training machines and issues short-term certificates for courses using the production CA. For their Grid schools and training events, however, they usually set up a local Grid rather than giving learners access to a central t-Infrastructure.

DemoGrid

DemoGrid [140] is a recent development from the Globus project which allows the creation of an instructional Grid environment deployed on local or cloud resources using Vagrant [141], KVM [142] or commercial cloud infrastructures such as Amazon Elastic Compute Cloud [143].

3.7.2 Other technology eLearning Practical Environments

In the wider area of technology education and training, there are a number of interesting t-Infrastructures not specific to Grid, but of great interest to this thesis because of their approach.

O'Reilly's Ellipse [144] is a learning plugin for the Eclipse open-source integrated development environment (IDE). The plugin includes features for communication between the learner and a tutor and for practical exercise submission directly through the IDE used to complete the practical exercises. Exercises are submitted to a tutor who manually marks them and there is no attempt made to automate this procedure as in the eLGrid, or of providing automated feedback.

O'Reilly's Head First Labs [145] has implemented an online component to accompany their educational books. Practical exercises are emailed to learners each day, and the author gives a feedback and Q&A session on the exercises each week. This does not include any submission facility, so learners work on the assignments themselves and ask questions in the Q&A session if they encounter problems.

The simulation environments discussed in section 3.4.3 are also relevant examples of eLearning practical environments.

3.8 Discussion

This chapter has presented the state of the relevant art in education, eLearning and Grid education. This section discusses the implications of all this for adaptive Grid eLearning and the XeL approach and attempts to set some guidelines for the design of an architecture for an XeL-enabled adaptive Grid eLearning system, and of an prototype implementation developed as an instrument to explore the questions of this thesis.

It is important that the instrument developed to explore the questions of this thesis should follow a sound pedagogic approach. The instrument and courses should follow the principles of constructivism, the dominant framework for eLearning systems today. The practical elements should be distributed throughout the course to take advantage of the distributed practice effect. The courses should facilitate active learning and include many opportunities to practice the skills taught. Worked examples should be included as well as practical exercises in order to reduce cognitive load. Both worked examples and practical exercises should be based on real-world examples in order to maintain learner interest and improve the completion rate of the courses and the practical exercises. Problem-based learning and other forms of active learning should also be considered and provided where appropriate.

The eLearning tools should provide opportunities to review and reflect on learning by providing feedback and showing progress through the materials. Feedback can be given on each learner's test answers and also on their practical exercises where these are integrated with the eLearning tools. Learners might be encouraged to review their progress by the simple expedient of providing an annotated menu which allows learners to see which of the course concepts they have viewed and how well they are considered to have understood them based on the results of their test questions and practical exercises.

The instrument and courses should, where possible, reuse and repurpose existing materials to create these pedagogically sound courses. A wealth of material already exists in repositories such as the ICEAGE Digital Library or the SuGI Grid Portal, as well as the materials on the websites of various Grid projects. While these resources tend to lack a pedagogical framework, they are still very useful materials produced by experts in the field. These should be reused where possible as learning content presented within the framework of a pedagogically-sound eLearning course. Where *multi-modal* materials are available, for example, where the same material is available in audio and text formats, then the effects of multi-modal delivery of material might also be explored.

The courses should support personalisation as this is a very successful method for overcoming some of the disadvantages of eLearning over traditional face-to-face learning. At a minimum, courses should adapt to the learner's prior-knowledge and their progress through the course. Adaptation to other attributes, in particular Grid-related attributes, should also be explored.

The instrument developed should be flexible enough to support remote self-paced learners and integration into face-to-face courses. It should be evaluated with both types of learning.

eLearning tools should be closely integrated with the Grid t-Infrastructure to allow seamless access to a practical execution environment (a core concept of XeL). The thesis should explore the effects of this, to determine if close placement of the practical exercises to the integrated practical environment can help to reduce cognitive load. The pedagogic implications should also be explored, with a specific focus on the possible affordances for personalisation and assessment of the integration of adaptive eLearning tools with a Grid t-Infrastructure. The eLearning and practical environments should be integrated, and the eLearning tools should use information about each learner's conduct of practical exercises (another core concept of XeL). The completion of practical exercises should be used by the personalisation engine as evidence of having learned particular concepts. This information should then be able to be used as pre-requisites for further concepts or courses, or as part of the assessment process.

Creating a true simulation of production Grid systems, or other large distributed computing systems, is very difficult due to the complex interrelations that exist between the various components and the network. While some work has been done on simulating distributed systems and algorithms [146] and Grids [147] the sheer size and complexity of a production Grid makes such an approach untenable. The approach proposed in this thesis, of using a fair replica of the real production Grid environment rather than a simulated environment, should be adopted to allow learners' practical exercises to be run in a realistic setting without the huge undertaking of building a simulated Grid environment.

Single sign-on capabilities of the Grid should be extended to the integrated eLearning tools. It should be possible for a learner to gain access to courses using their Grid credentials, and preferably, this should be extended to appropriate federated identity credentials.

Broader aspects relating to the idea of a personal/group XeL work-bench should be explored. This thesis should look at whether the eLearning tool can form a useful interface to the real production Grid and whether the learners respond positively to this. The reusability of the tools, and in particular the practical exercises, by learners for continuing professional development might also be considered. The potential for collaborative learning should be also be examined.

Chapter 4

XeL Architecture

4.1 Introduction

This thesis sets out to investigate the potential for using adaptive eLearning technologies to enhance Grid education, in particular the question of whether a learner's practical exercises running on the Grid can feed back into the personalisation of adaptive Grid eLearning courses, and whether this gives any benefit to the learner. That is, whether the concept of adaptive executable elearning, or *XeL*, is feasible and beneficial in Grid eLearning.

In order to explore this question, the concept of XeL must be more clearly defined and the main components of an XeL system identified. This chapter describes XeL in more detail and suggests a possible XeL architecture. It goes on to briefly describe the prototype implementation of this XeL architecture developed for this thesis.

4.2 XeL Overview

Existing Adaptive eLearning systems allow personalisation on a range of learner attributes, as described previously in Section 3.4.2. This thesis proposes the concept of XeL, adaptive executable eLearning, whereby the learner's practical exercises become another attribute on which personalisation can occur.

Section 2.4 outlined the potential attributes of an XeL eLearning system:

1. Distributed pedagogically-sound self-paced adaptive Grid education via eLearning
2. Support for integration into face-to-face courses
3. Integration of t-Infrastructure for practical exercises within the eLearning environment
4. Capture of relevant information about practical exercises by the eLearning tool
5. Leveraging of single sign-on capabilities of the Grid
6. Adaptation to the information about learner's conduct of practical exercises
7. Ability to treat satisfactory completion of practical exercises as pre-requisites and for assessment purposes
8. Ability to revisit adaptive eLearning including integrated practical exercises
9. Possible extension to become an integrated problem-solving "personal/group XeL work-bench"

Among these, the core concepts of XeL are items (3), (4) and (6). These are the integration of the t-Infrastructure for practical exercises within the eLearning environment, and the ability to both capture and use information about the learner's conduct in the t-Infrastructure for the purposes of course personalisation.

Any XeL-enabled eLearning system must therefore include not only adaptive course presentation software, but also a t-Infrastructure. The t-Infrastructure should not be a separate entity, but rather it must be integrated with the eLearning system, with channels of communication between the two to allow transfer of information about learner activities.

The integrated t-Infrastructure should not be a canned simulation, but a real system or replica of same, so that learners can fully explore the functionality of the Grid software which they will use.

Information about the learner's practical exercises run within the t-Infastructure must be made available to the eLearning system. This information should include not only the status of the exercises (whether the learner has completed an exercise or not), but also the output and results of the exercises, as this is the best way to allow meaningful adaptation to the activities of the learner.

4.3 Architecture of an XeL-enabled eLearning system

Adaptive eLearning systems allow personalisation of courses based on various attributes of, at a minimum, the learner and the subject being learned. These attributes are stored in models which are used by the adaptive eLearning system to create a personalised course for each learner.

An adaptive eLearning system typically consists of a number of basic components:

- a set of models which store information about the subject domain, learner and other relevant entities
- an adaptive eLearning engine which uses these models and other pertinent information to produce a course
- optional authoring tools

Interbook [148], for example, consists of an authoring tool which generates html pages from a rich text format (RTF) file containing the course content which is produced by the tutor in a particular format. The InterBook server then generates a personalised menu to navigate these html pages at run-time, based on a *domain model* and *student model*.

AHA! [149] has a similar architecture but it combines the Domain Model with an Adaptation Model, which allows for more flexibility in course personalisation.

The APeLS [64] multi-model approach forces a further separation between the generation and presentation of the course, as each concept has potentially several different "candidate resources", or html pagelets which could be presented to fulfil the learning objectives of a particular concept. The domain model must therefore not be tied to individual learning content resources. The choice of which candidate resource to select is made by the adaptive eLearning engine based on the interaction of the various models. The APeLS architecture introduces a web-based front-end course presentation application (or a course renderer in the terminology of APeLS) component to the architecture, which is a separate component from the models and the adaptive engine.

Some systems combine one or more of these components, but the most common configuration is usually something akin to this, although the language used to describe them and the exact details of architecture vary among implementations.

An XeL-enabled system must add a number of additional components:

- a t-Infrastructure which provides an environment for learners to try out what they learn
- one or more XeL modules which feed information about learner's conduct in the t-Infrastructure back to the adaptive eLearning engine
- additionally, because of the nature of the t-Infrastructure as a realistic replica of a production infrastructure, it is usually not possible to embed it directly into the eLearning application as one might a simulation. This means that a networking layer may also be required which provides for communication between the adaptive eLearning engine and web-based front-end course presentation application, and the t-Infrastructure
- again, because of the complexity of creating and maintaining a realistic t-Infrastructure, additional management tools may be required

An XeL-enabled prototype implementation called eLGrid was developed in order to explore the questions of this thesis. It follows the architecture proposed here and consists of the following components:

(a) a training infrastructure (t-Infrastructure)

(b) a front-end web-base eLearning system (eLGrid eLearning application)

(c) an adaptive engine which creates personalised courses

(d) a set of XeL modules to collect information about learner's conduct of practical exercises and feed these into the adaptive engine

(e) metadata models which the adaptive engine will use for personalisation

(f) a networking layer which isolates the t-Infrastructure from the production network while allowing communication between the t-Infrastructure and the adaptive eLearning application and web-based front-end application

(g) a set of management services, including authoring tools and system administrator tools to simplify configuration and maintenance of the system

(h) various external services which may be required for operation of the system

(i) a body of courses to attract learners to the system in order to help in its evaluation

This architecture is illustrated in Figure 4-1.

This chapter describes the components of this architecture in more detail. It also summarises the design decisions taken in developing a prototype implementation of the XeL architecture, both to support experimental investigation of the validity of the XeL concepts, and as an exploration of those issues of feasibility which must be considered in order to fulfil the aims of this thesis.

4.4 The t-Infrastructure

This section describes the t-Infrastructure design requirements and decisions.

4.4.1 t-Infrastructure Design Requirements

The generic requirements regarding t-Infrastructure have already been discussed in section 3.8, viz, to closely integrate the eLearning tools and the t-Infrastructure, to use a fair replica of the production Grid environment, and to enable single sign-on capabilities. Consequential requirements are discussed below.

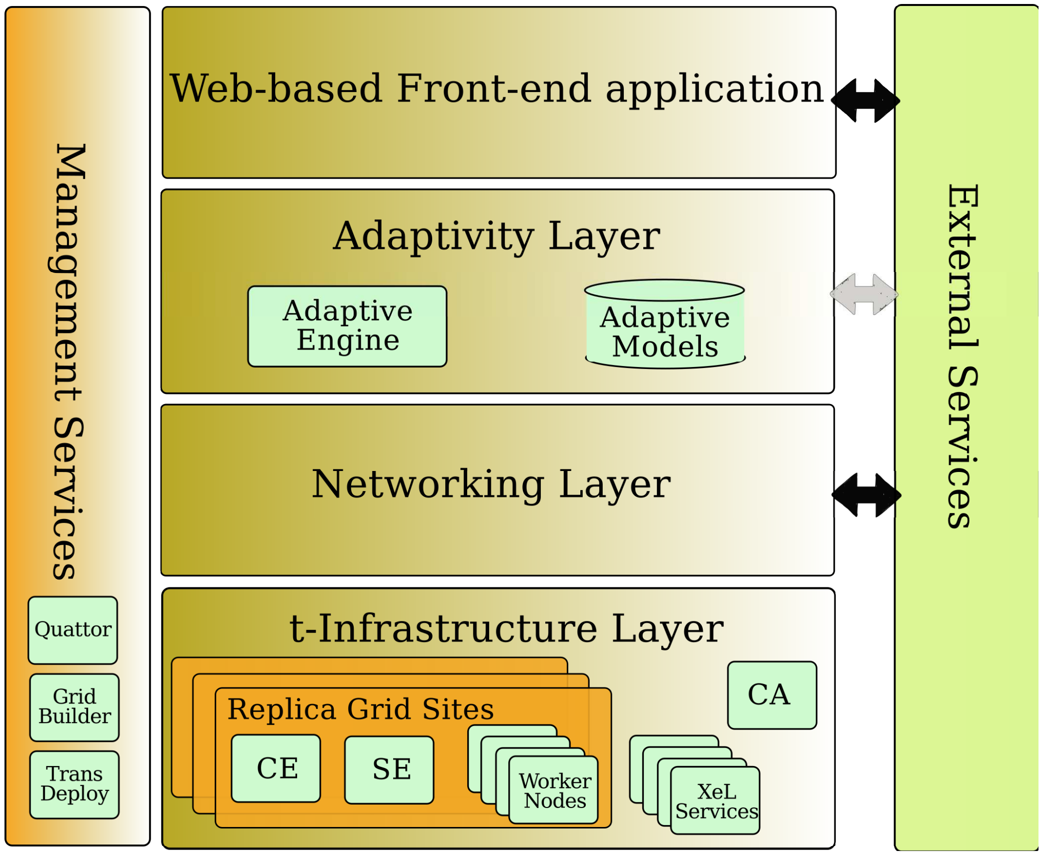

Figure 4-1: A proposed XeL architecture

Integrate eLearning and Grid environment

The benefits of integration were already discussed in Sections 2.2.3 and 3.8.

Simulate the production infrastructure

Training users on an infrastructure which differs substantially from the live production environment may prove counter-productive as these users may simply end up frustrated if they find that certain services and middleware versions with which they trained are not available on the production infrastructure. The more closely the learning environment can mimic the production infrastructure, the shorter the learning curve for users moving from training into production use.

For this reason, some projects have chosen to train users on the production infrastructure. Issues such as poor response times and security risks (mentioned later in this section), however, make this a potentially unsatisfactory solution, and many of the projects who have taken this approach in the past have since begun to use the dedicated resources of the GILDA t-Infrastructure. For example, the VOCE [138] project, now devolved into the National Grid Infrastructures of the Central European region, no longer maintains a training VO.

Enable single sign-on

Again, see Sections 2.2.3 and 3.8.

Dedicate resources to guarantee QoS

Grid jobs that are run as part of a face-to-face or online training course must be responsive as the learner may have to wait for the results of one job before progressing to the next. If the learner's short training jobs are queued behind longer production jobs this can give the impression of unreliability.

Different t-Infrastructures tackle this problem in different ways, GRNET use high priority queues for all training jobs. The VOCE and GILDA approach of dedicating specific resources for training is a more reliable way to ensure fast responses.

In addition to the problems for learners outlined above, the bursty nature of traffic in a face-to-face training session, where multiple learners may submit a large number of jobs at the same time, has the potential to disrupt production jobs.

Furthermore when training Systems Administrators it may in fact be necessary to 'break' part of the Grid in order to teach administrators how to resolve common problems. This is difficult to do if the resources are not independent from the production Grid as the misconfigured component(s) could affect the Quality of Service (QoS) of production jobs. Any attempt to train for scenarios where issues arise simultaneously will only exacerbate the impact on the QoS of production jobs.

For these reasons it is desirable to use dedicated resources for training activities where possible. This will improve the QoS for both the learners and other Grid users. The burden of managing separate resources, however, militates against a fully independent set of dedicated resources, and so the current trend is towards co-location and some form of marriage of production and eLearning resources [150].

Isolate training testbed for security

It is desirable to train as many users as possible in Grid technologies, but giving access to a large number of potentially unknown and untrustworthy users has its own attendant problems.

When inexperienced users are given access to the Grid it is always possible that they will inadvertently cause some problem. In addition to the danger of learner mistakes, the possibility of malicious action by learners cannot be ignored.

As the target learner community broadens from HEP to the wider Scientific community and then to industry, these dangers are increased.

Make easy to use and configure

The integrated t-Infrastructure must be very easy to use for both learners and administrators.

From the learners' point of view they do not wish to become proficient in using a complex eLearning system before they can learn about the Grid.

It is not simply the learners, however, for whom the system must be easy-to-use. The systems administrators and tutors will need to be able to set up new sites for training courses and to handle ongoing updates and maintenance of the existing sites. This should be made as simple as possible as few organisations have significant resources to dedicate to training activities.

If the systems administrators have a fabric management system for installation and configuration of all of their production resources then ideally the t-Infrastructure should also use this system. If this is not done then there is an additional workload involved in managing the t-Infrastructure machines. The systems administrators will naturally give higher priority to the production infrastructure. Thus having a separate t-Infrastructure management system can be self-defeating, in that it can lead to the t-Infrastructure being ignored in favour of the production infrastructure.

4.4.2 Choice of t-Infrastructure

The t-Infrastructure implemented for this thesis is a replica of the Grid-Ireland production infrastructure. It uses virtualisation to allow a large number of replica Grid sites to be hosted on a relatively modest number of physical machines. The details of the t-Infrastructure implementation are described in Appendix A.

As mentioned in section 3.7.1, the GILDA t-Infrastructure is a widely used Grid t-Infrastructure used by several EU Grid projects. GILDA could have been adopted as the t-Infrastructure for an XeL-enabled Grid eLearning system, however, the prototype implementation for this thesis uses a custom t-Infrastructure which is based on virtualisation. This decision was taken for two main reasons.

Firstly, the requirements above call for a close replica of the production Grid system, and it was decided to replicate the Irish production Grid infrastructure, as Irish users were expected to be the primary users, at least initially. Replicating a full production Grid infrastructure on real resources is not feasible, and in order to create a realistic replica while using minimal hardware it was necessary to use virtualisation. The Irish Grid system already had a virtualised testing infrastructure and it was decided to use a similar configuration for the t-Infrastructure.

Secondly, the use of a custom t-Infrastructure allowed more control over the exact configuration and made it easier to experiment with methods of integration between the t-Infrastructure and the eLearning application, and to investigate possible ways to implement the XeL modules.

Over the lifetime of this thesis, GILDA has moved in the direction of providing virtual hosts for Grid nodes and would thus now be a more attractive option. In addition, the initial implementation has shown that it should be possible to implement the XeL modules and integration with the eLearning application on another t-Infrastructure without having complete control of the configuration of that infrastructure. Therefore the option of using GILDA as a t-Infrastructure may be revisited in future.

4.4.3 Certificate Authority

Grid authentication is achieved through the use of X.509 digital certificates. These are issued by a Certificate Authority (CA). An XeL-enabled Grid eLearning system must have the ability to issue Grid credentials to its learners, either by connecting to an external CA service or by implementing its own CA within the t-Infrastructure.

While either solution would be acceptable, the approach taken in the prototype implementation for this thesis was to create a dedicated XeL CA which issues short-term limited credentials valid only on the t-Infrastructure. This allows the XeL administrators to have more control over the duration and scope of the credentials issued and seemed a more prudent approach, at least in these early stages. The implementation details of the XeL CA are described in Appendix A.

4.5 The Front-end eLearning Application

This section looks at the design requirements and decisions relating to the front-end eLearning application.

The generic design requirements have already been discussed in section 3.8, viz, to follow the principles of constructivism and support active learning via pedagogically sound courses. These and some consequential requirements are highlighted below.

- *Principles of Constructivism*: See Chapter 3.
- *Active Learning*: See Chapter 3. In particular, an annotated menu system that allows learners to review their progress should be explored.
- *Pedagogically-Sound Courses*: See Chapter 3. All materials reused from other sources must be presented within a sound pedagogic framework, and adherence to relevant standards should be explored where appropriate in order to facilitate the reuse and sharing of materials.
- *Support personalisation of courses*: See Chapter 3. In particular, an annotated menu system that provides navigational guidance to the learner, without constraining their actions, should be explored.

- *Minimise Extraneous Cognitive Load*: See Chapter 3.

- *Take advantage of the distributed practice effect*: See Chapter 3.

- *Provide timely feedback*: Timely feedback should be provided on the practical exercises. Furthermore, the integration of the t-Infrastructure and practical exercises may allow the results of learner's practical exercises to be used by the eLearning application in various ways, for example, in creating subsequent learning content and examples, or by presenting the results in the form of a learning journal or e-Portfolio. The pedagogic implications of these affordances of integration should be explored.

- *Ease-of-use*: Finally ease-of-use should be a primary consideration throughout the creation of these tools. Grid technologies are often complicated and difficult to learn and the design of the eLearning application must not add to this complexity.

4.5.1 User-Interface Design

Many Grid projects don't worry too much about user-interface or usability issues, simply providing APIs or command-line interfaces which are expected to be used by experienced UNIX-savvy physicists and programmers. This is changing, however, and first-generation user-friendly graphical interfaces (such as GANGA [151] and the Migrating Desktop [152]) and workflow portals such as P-Grade [153] are now maturing, and even more advanced graphical interfaces (such as WS-PGrade) are gradually being introduced.

Any XeL-enabled front-end application must act as an interface of sorts to the Grid, insofar as to make it possible to launch Grid applications and run practical exercises from within the eLearning application environment. Unlike traditional Grid user interfaces, however, an eLearning application is aimed specifically at inexperienced users. These users are unlikely to have had significant exposure to the Grid.

It also makes sense to assume that a proportion of learners will have little experience with UNIX or Linux, programming and computing in general. Indeed, this proportion should increase over time as Grid penetration grows beyond the traditional, technically-competent and UNIX-savvy demographic of the current heavy Grid user communities such as high-energy physics, life sciences, astronomy and astrophysics, earth sciences and computational chemistry [154, 155].

It is therefore vital to take into account user-interface issues and design the interface to be easy to use and navigate. Many of the Grid concepts being introduced have a high intrinsic cognitive load, that is, they are relatively complicated and difficult to learn; the design of the system must minimise extraneous cognitive load so as not become an additional barrier to learning.

eLearning courses tend to experience high drop out rates [156] compared with traditional face-to-face courses, and learner satisfaction with the eLearning system is a key indicator in their decision to drop out [156]. Learner satisfaction is clearly linked to ease-of-use [157] and thus a careful approach to user-interface design may improve the probability of learners completing courses.

Furthermore, studies have shown how poor navigation design can adversely affect learning itself [158]. Similarly page-layout which does not conform to a learner's expectations can increase cognitive load thus adversely affecting learning [159].

Therefore the degree of success in designing a user-friendly interface becomes a crucial factor in determining the ability to get results from this thesis. The prototype implementation developed for this thesis aims to follow accepted guidelines in designing web-based applications in order to ensure usability. More detail on this is presented in Appendix A.2.

4.5.2 Authentication

The XeL eLearning application and t-Infrastructure should be accessible to learners with a Grid certificate, but ideally it would also be possible to request an account and log in without, or before being issued with, Grid credentials. One of the courses published via the prototype implementation developed for this thesis takes the learner through the process of applying for a Grid certificate, so it makes sense for learners without a certificate to have access so that they can complete this course. Once completed, learners should be able to use their Grid certificate for subsequent authentication.

The use of Grid certificates for authentication means that the XeL front-end application knows the learner's DN and is be able to use this to obtain information about the learner's practical jobs.

The implementation details of the authentication mechanisms of the prototype implementation developed for this thesis is given in Appendix F.

4.5.3 Adaptivity

The main method of adaptivity implemented in the prototype system is adaptive navigation.

Once authenticated with the eLearning application, the learner is presented with a course list, and upon choosing a course there are two possible ways for learners to navigate through the course concepts; *Next* and *Previous* buttons and a menu annotated using a *Traffic-light* metaphor.

Next and Previous buttons - Direct Guidance

The *Next* and *Previous* buttons in the eLGrid eLearning system use the principle of Direct Guidance [14] to help the learner navigate through the course content.

The *Next* button recommends the best resource for the learner to view next. This may not always be the next item in the menu. If the learner has already viewed and passed a test in the next concept, for example, then the *Next* button will skip over that concept and bring the learner straight to the next untested concept. Similarly the *Previous* button skips over tested concepts. These tested concepts are still available via the menu, however, should the learner wish to view them again.

Menu - Link Annotation and the Traffic-light metaphor

A course menu is displayed on the left-hand side of the page. The menu is created adaptively and uses link annotation with a traffic-light metaphor [14].

Menu items have a coloured icon placed beside them which provides the user with additional information about the concept linked to each item.

(a) A green icon indicates that the learner is ready to view a concept. It is also used beside course sections to indicate that the learner can view that section.

(b) An amber icon indicates that the learner has already viewed a concept and that they are now ready to be tested in this concept if they so choose. If no test exists in the course then the amber icon will not be used, instead the concept will be marked with a lilac icon (see (d) below) once it has been viewed. If a test is later added then the traffic light colour will change the next time the learner visits the course.

(c) A red icon indicates that the learner has not yet viewed or been tested for the prerequisite materials for this concept and is thus probably not ready to view this material.

(d) A lilac icon indicates that the learner has been successfully tested for this concept, or that they have viewed the material if no test exists for the concept.

(e) A grey icon indicates that the associated course section is not a necessary part of the course for that learner. It might be used in a case where a course is used for students from different disciplines, for example, if a particular section is only required by some of these students.

The eLGrid front-end is shown in Figure 4-2, where the *traffic-light* annotated menu is visible on the left of the screen.

4.6 The Adaptive Engine

An Adaptive eLearning Engine creates personalised courses at run time by adapting the content or the navigation based on the interaction of a number of *models* such as *Pedagogic Activity Sequence* (narrative), *Subject Area* (concept space), *Candidate Learning Resources* (learning content), Learner Model, etc.

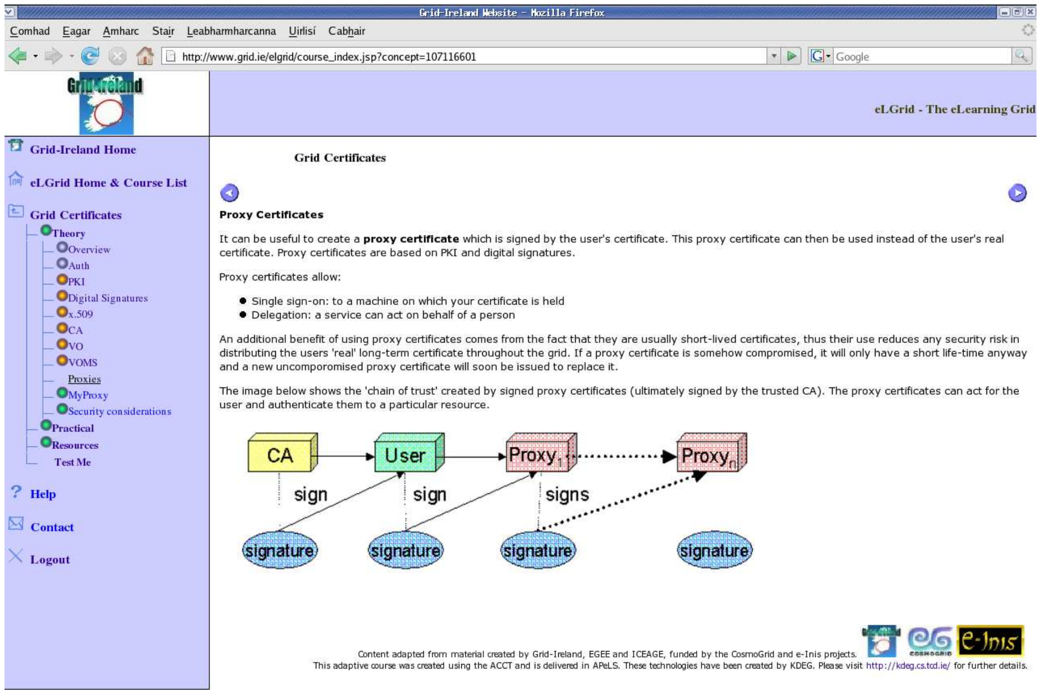

Figure 4-2: The eLGrid front-end

The adaptive engine selected for the prototype implementation for this thesis is the Adaptive Personalised eLearning Service (APeLS) developed by the Knowledge and Data Engineering Group in Trinity College [64]. This tool uses a multi-model approach which allows flexible and extensible course personalisation. More details are given in Appendix A.

4.7 XeL Modules

A common problem of eLearning is to assess whether the user has really understood the material in the absence of a human tutor who can best make this assessment. Traditionally it has been difficult to employ more than multiple-choice assessment without the intervention of a human tutor.

As discussed in section 3.4.3, pre-recorded simulations with interactivity can allow learners to practice tasks which can be assessed by the eLearning or simulation software. However, such simulations are necessarily limited.

Traditional lab sessions ask learners to use the software tools themselves to carry out this practice, but these also require a tutor to perform the assessment.

An ideal solution would be a way to use the real tools, but somehow to feed information about the learner's conduct in these tools back to the eLearning system for use in assessment and personalisation.

To achieve this the tools would require some instrumentation, ideally a framework for instrumentation of tools to allow them to be used in this fashion.

XeL allows the instructions to be presented on screen, in text, screenshot, video, or other form, then the learner performs the tasks in a t-Infrastructure using the real software. The results are then fed back to the eLearning system via plugin modules in order for the results to be used for assessment or further personalisation of the course

An XeL-enabled adaptive eLearning system uses e-Infrastructures as tools to help train scientists in the use of said e-Infrastructures, and this presents a fruitful opportunity to apply leading eLearning principles such as closed-loop eLearning.

The users' learning tasks (job submissions, etc) actually run in the same environment in which the eLearning software runs. As a consequence, the results are available to the eLearning software for use as inputs to the adaptivity engine and can inform subsequent instruction or exercises. These types of adaptivity require adaptation plugins which can be written to refer to the Grid information system or the execution results of jobs which the learner has submitted.

Calling this concept "Executable eLearning" (XeL) indicates the integration of the eLearning application and the execution environment, and the central role of this execution environment in the course adaptivity.

There are two proposed initial scenarios or use-cases for XeL: capturing the job status and capturing the job output. These are discussed below in sections 4.7.1 and 4.7.2. Both result in updating the learner model to allow the adaptive personalisation of new content.

In addition, because XeL relies on an integrated training infrastructure (t-Infrastructure), it could allow dynamic adaptation based on aspects of this infrastructure such as the available software versions or libraries, or even the types of middleware or Grid applications which are available to the user. This requires further integration with the t-Infrastructure. The content metadata must store the requirements in terms of software, middleware, applications, etc.

4.7.1 Capturing Job Status

The first scenario involves the result status of learners' jobs being fed back into the learners' profiles to be used in course personalisation.

An XeL service could be notified when a learner job is submitted to to the t-Infrastructure and will then poll the Grid Information system for the status of the job. Alternatively, the XeL module could be bundled and submitted with the learner's job in some way, thus avioding the need to install anything on the t-Infrastructure.

If the job completes successfully the related concept can be marked as understood in the learner's profile. The XeL service will send a notification to the eLGrid application which will perform the actual database update.

The profile update will have the same effect as if the learner had passed a test for that concept, so it will be possible to use both tests and practical exercises in this way to update aspects such as the learner's navigation path through the course.

Because the Grid provides generic services such as the Information Service where all information about a job can be published, it should be possible to handle a wide range of different job types in this way, although the information available via the Information System is rather limited.

By bundling the XeL module with the learner's job the limitations of the Information System are bypassed, but the learner may need to perform some extra steps in order to enable XeL.

If using other interfaces to the Grid, such as workflow portals, etc. these may provide their own interfaces and APIs for the XeL modules to get the job status.

4.7.2 Capturing Job Output

The second scenario involves the output of the learner's jobs being captured and used to generate instructional content, examples or further exercises.

In this case the output files must be copied back to the eLGrid application server where they can be incorporated into newly generated learning resources and content.

XeL can support more complicated adaptivity beyond the adaptive navigation. By uploading and storing the results of the job, the result can be checked and instant feedback provided to the learner. Summative and formative assessment is possible in this way. These can be made available to the student so that they can later review their results and progress in a form of learning journal or e-portfolio. Results can further be used to adapt the content used in subsequent examples or to create tailored worked examples or exercises.

As an example of what is involved, in the eLGrid *WebCom* course learners are asked to complete a number of simple practical exercises and finally to complete a larger exercise which includes all aspects of the previous exercises: recursion, creating custom nodes, control-driven and data-driven computing models.

Learners must create a custom node to download an input value N from the eLGrid web server. They then create a recursive WebCom Graph to calculate the *Nth* Fibonacci number. The result must be passed to a special eLGrid node which is provided to them. This node uploads the result to the eLGrid server where it is checked and stored, and a response code returned to the eLGrid WebCom node to indicate whether the answer is correct or incorrect. The eLGrid WebCom node then outputs a string to the user which includes their response along with a message indicating whether they were correct or not. When the correct response is received at the eLGrid server, the learner's profile is updated to indicate that they have successfully learned the related concept.

At the server end, a handler is required in order to process the job output. This might simply store the result, or it may check it for correctness and provide feedback. It might also perform more complex post-processing of the data in order to generate new content.

4.8 Metadata Models

Adaptive courses are typically created at run-time based on metadata models associated with the learners, content and other aspects of the course. The tools selected for the prototype implementation for this thesis allow a certain amount of flexibility in the exact format of the models. This enabled new metadata formats to be created which, while based on existing standards, contained additional information specifically relevant to Grid eLearning.

4.8.1 Learner Model

The learner model in an adaptive eLearning course stores information about the learner, their background and competencies, which can be used by the adaptive engine to create a personalised course.

A typical learner model might include the learner ID, name, and other basic information, along with details of their learning preferences and a list of concepts which they have completed. The APeLS learner model is extensible; new XML entities can be added and the appropriate XSLT transforms created to take advantage of these.

For the purposes of this thesis two new attributes were added, the learner's *domain* (Computer Science, Physics, Bioinformatics, etc.) and their *Virtual Organisation*. Personalised courses can thus be delivered which present different content to the learner depending on their background or the experiments or Virtual Organisations of which they are members.

In the prototype implementation developed for this thesis, the learner's domain is used to adaptively present certain information. Link annotation is used to change the colour of the link icons in accordance with a traffic-light metaphor. Some course sections may only be required by learners from a particular domain or VO, and these can thus be adaptively presented.

Ideally these would be automatically pulled out of the learner's Grid credentials in order to populate the learner profile, however, for the prototype implementation these were generally left blank as many users logged in using user name and password rather than their Grid certificate. In future, as more users authenticate with certificates, more use could be made of this feature.

An example of a learner model metadata file for the prototype implementation of this thesis is given in Appendix D.

A web-based application was developed to create and populate an initial learner model when a new eLGrid user account is created. This initial learner model is essentially a blank profile with default values, the application does not attempt to find any information about the learner at this stage.

4.8.2 Content Model

For existing APeLS courses KDEG have used a schema for their content metadata which is not based on any particular standard but which was designed as a canonical form. KDEG have recently expressed a wish to move towards standards based metadata.

An example of a metadata file in the format used by some of the APeLS courses in the past is shown below. An example of the new format used by eLGrid is shown in section E.

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>

<pagelet>

<general>
  <identifier>gma</identifier>
  <title>Grid Monitoring Architecture GMA</title>
  <keyword>GMA</keyword>
  <description>Introduce GMA framework</description>
  <language>en</language>
</general>

<pedagogical>
  <objectivestaught>
```

```
    <objective>GMA</objective>
  </objectivestaught>
  <supportedlearningstyle></supportedlearningstyle>
  <semanticdensity></semanticdensity>
  <displayarea></displayarea>
</pedagogical>

<technical>
  <location>
    http://blacktower.cs.tcd.ie/rgma_course/vdb-gma.html
  </location>
  <format></format>
  <requirements>none</requirements>
  <size>817</size>
</technical>

</pagelet>
```

The potential rich adaptivity in eLGrid courses requires more complex metadata than the example above. For example, it may be necessary to store information about the VO which published the content, or the discipline (e.g. physics, chemistry) for which the material is suitable, etc.

The author has thus created a format for a Grid content model based on the IEEE Learning Object Metadata (LOM) [160] [161] standard, with Grid-specific extensions. Some preliminary discussions were concluded with members from the National e-Science Centre (NeSC) in Edinburgh who developed the ICEAGE Grid People Registry and the ICEAGE Digital Library in order to ensure consistency in the vocabularies used within this format. An example of this new content model format is given in Appendix E.

IEEE LOM is not a stand-alone standard, rather it "is intended to be referenced by other standards that define the implementation descriptions of the data schema so that a metadata instance for a learning object can be used by a learning technology system to manage, locate, evaluate, or exchange learning objects" [160].

Thus for eLGrid it is used in conjunction with other standards which have further defined and clarified certain areas of the LOM specification. These standards are IMS Meta-data Best Practice Guide for IEEE 1484.12.1-2002 Standard for Learning Object Metadata [162] and the Cancore [163] and UK LOM Core [164] Application Profiles.

Furthermore, the IEEE LOM standard is a base schema, designed to be extended to incorporate new developments in eLearning including adaptivity [160].

Thus where it is necessary new fields have been added to the eLGrid metadata schema to facilitate the desired types of adaptivity. This has been done with reference to the guidelines in the IMS best-practice document, for example, wherever a field is repurposed an additional instance of that field with a different vocabulary source is added, rather than replacing the field entirely. This allows the creation of a rich metadata schema which supports complex adaptivity and XeL but which is also standards compliant and which satisfies the minimum requirements for interoperability and metadata sharing.

Many fields support a limited value space which must be defined by a vocabulary. While LOMv1.0 defines many of these vocabularies within the standard not all vocabularies are sufficiently rich for the purposes of eLGrid. It is possible to use other vocabularies in addition to the LOM ones simply by adding additional instances of the element in question with a ”source” field which identifies the vocabulary. Where possible other standard, well-recognised and mature vocabularies have been used, but for some metadata elements it was necessary to create custom vocabularies. This was particularly true for Grid-specific metadata elements where mature, standardised vocabularies do not yet exist.

The OGSA glossary [165] and the OGF glossary [166] have been used as a source for many Grid-related terms and only where there was no suitable alternative has an entirely new vocabulary been created. Where this has been necessary, an effort has also been made to drive forward the development and standardisation of these vocabularies through the OGF Education and Training Community Group [4], the ICEAGE project [2] and other fora.

The sections below describe some of the possible entities which should be covered by a Grid eLearning content model, and Section E shows a sample eLGrid content model file.

Operating System and other platform requirements

XeL involves actually running learners' practical exercises in an integrated t-Infrastructure. Thus it is useful for eLGrid to know something about the t-Infrastructure requirements of a practical exercise.

The LOM standard allows the intended operating system for the eLearning materials to be specified in the *technical.requirement.type.value = "Operating System"* element. This could potentially be repurposed to specify the required OS for the practical exercises. For this purpose it would be necessary to extend the vocabulary beyond the values listed in the UK LOM specification

- ms-w33
- ms-windows
- macos
- unix
- multi-os
- none

Various flavours of Linux or other UNIX systems may need to be added as some Grid software only runs on certain distributions.

For the time being, however, the eLGrid content model simply uses *"none"* in this field.

The *otherplatformrequirements* field in the specification is a text field allowing the hardware and software requirements for XeL practical exercises to be specified more flexibly and thus this field is used.

Educational Level

The LOMv1.0 controlled vocabulary allows for the educational level of the learner to be specified, but the highest level is *Higher Education*. As many real or potential eLGrid users are postgraduates, postdoctoral researchers, academics or in industry, this classification is not particularly suitable.

```
<context>
  <source>
    <langstring xml:lang="x-none">LOMv1.0</langstring>
  </source>
  <value>
    <langstring xml:lang="x-none">Higher Education</langstring>
  </value>
</context>
```

The current version of the eLGrid content model generally uses "Higher Education" here, but a new vocabulary should be developed to cover the different levels of higher education and beyond into continual professional development.

The results of the experiments carried out as part of this thesis which seem to suggest differences in the usage of the system between learners at different levels suggest that this could be a very useful modification to the content model, allowing personalisation of courses for undergraduates, postgraduates, postdoctoral researchers and academics.

VO

Some scenarios require the publisher's or learner's VO (or both) to be identified. For example, some educational material may be adaptively published by a VO to members of that VO only. This would require the publishing VO to be stored in the content metadata so that it can be compared with the learner's VO which will be extracted from the digital certificate at login time and stored in the APeLS Learner Model.

The following section of the metadata indicates the publishing VO (in this case the Irish CosmoGrid VO):

```
<classification>
  <purpose>
    <source>
      <langstring xml:lang="x-none">
        http://www.cs.tcd.ie/Kathryn.Cassidy/taxonomy.html
      </langstring>
    </source>
    <value>
      <langstring xml:lang="x-none">VO</langstring>
    </value>
  </purpose>
  <taxonpath>
    <source>
      <langstring xml:lang="x-none">
        http://www.cs.tcd.ie/Kathryn.Cassidy/taxonomy.html
      </langstring>
    </source>
    <taxon>
      <entry>
        <langstring xml:lang="en">cosmo</langstring>
      </entry>
    </taxon>
  </taxonpath>
</classification>
```

As indicated above some publishers (e.g. VOs) may put restrictions on the use of their educational materials. In order to have a flexible, extensible metadata schema, any such restrictions must be able to be stored in the metadata.

One example would be the VO-specific material outlined above. The type of restriction must be stored in the content metadata in order to distinguish between restricted materials published by that VO and other materials which may not have such restrictions.

The following entities are used to indicate usage restrictions placed on the material by the publishing VO.

```
<copyrightandotherrestrictions>
```

```
  <source>
    <langstring xml:lang="x-none">LOMv1.0</langstring>
  </source>
  <value>
    <langstring xml:lang="x-none">yes</langstring>
  </value>
</copyrightandotherrestrictions>
<description>
  <langstring xml:lang="en">VO internal</langstring>
</description>
```

This indicates that the resource can only be used by members of the publishing VO.

Learners may have different roles within a VO and it may be necessary to present different materials to, say, a project scientist than to a system administrator. The role for which the materials are best suited should be stored in the metadata, where different materials might be created for different learner roles.

Vocabularies

While various controlled vocabularies are available for use with the LOM standard, there are some areas where these were found to be lacking.

Disciplines Much of the content will be the same for any type of learner, but some content, particularly exercises, examples, case studies, etc., may need to vary based on the discipline of the learner. For example, an MPI-programming assigment could be tailored to the learner with physics students being asked to solve a heat dissipation problem while an environmental scientist is given a parallel population dynamics simulation.

Regarding the discipline or disciplines for which the learning content is suited, this information should be stored in the content metadata. Regarding the discipline of the user, in a non-Grid-enabled adaptive eLearning system this might be determined at registration time or via a pre-course learner questionnaire; for a Grid-enabled system the Grid user credentials generally contain some information about the user's organisation which might be used to identify their discipline, so alternatively, the user's VO could also be used to make an inference.

It is necessary to define a set of Disciplines to identify the suitable audiences for learning content. Such a vocabulary might include entries such as:

- *general* a catch-all for learners not belonging to any specific discipline
- *physics* for physicists
- *biochem*
- *lifescience*
- *computational*

Some classifications for such disciplines, or more properly, for subject areas do already exist, for example Dewey [167] and LCSH [168]. However these are very heavy classifications for the purposes of Grid education.

It is instructive to look at the approaches of other Grid projects which require discipline classification. For example the ICEAGE [2] digital library [94], rather than using an existing taxonomy, has chosen to use the Web2.0 folksonomy [169] approach.

eLGrid might also take this approach, or define a 'loose' taxonomy based on the user-community types identified by existing Grid projects.

OGSA Glossary Vocabularies are required for Grid-specific terms or items which do not already exist as formal taxonomies or vocabularies. The OGSA Glossary has provided many of these terms.

Content Types A number of different types of content must be supported and there may be different metadata requirements for each. Types of content include: text/html, images, video, audio and animations.

4.9 Networking Layer

The importance of ensuring that the t-Infrastructure mirrors the production infrastructure as closely as possible has already been discussed in Section 4.4.1, however, an additional design requirement set out in that section was the need to isolate the t-Infrastructure in order to protect the production services from accidental or malicious damage. In order to do this a networking layer which isolates the t-Infrastructure from production, while still ensuring that the services within it are accessible to learners and that they adequately mirror the production services, is a vital component of any XeL-enabled infrastructure.

In the prototype implementation for this thesis, mirroring of the production infrastructure involved using the *real* production IP addresses and network configurations for the t-Infrastructure machines, and the isolation from the production network was achieved by a firewall using a combination of NAT, iptables and arptables. The details of this implementation are described in Appendix B.

4.10 Management Services

A number of management services are required to ensure the smooth operation of an XeL-enabled infrastructure. These may include tools for installation and configuration of nodes, monitoring, etc., but also include a set of tools for course creators and tutors to ensure that they can create and modify courses, view learner progress, and perform other necessary functions.

In the prototype implementation for this thesis the management tools for system administrators consist of a number of services largely developed by the Computer Architecture and Grid Research Group at Trinity College Dublin [170], such as the GridBuilder [171] and TransDeploy [172] tools.

The course administration tools include the Adaptive Course Construction Toolkit (ACCT) [173] developed by Dr. Declan Dagger of the Knowledge and Data Engineering Research Group in Trinity College, along with a number of tools created by the author in order to further simplify creation and management of courses and learners. These tools are described in more detail in Appendix C.

4.11 External Services

An important focus in the design of this architecture was to enable reuse and repurposing of existing materials and resources, and to integrate with existing Grid services. The possibility of integrating third-party Grid training services and resources was also investigated and some limited integration was implemented.

4.11.1 Digital Library

The ICEAGE Digital Library [94] is a digital repository containing a huge number of educational resources including presentations, audio recordings of lectures, books, podcasts, tutorials, videos, exercises and assessments. These materials are freely available under the Creative Commons licence [117]. Educators can reuse and modify them for their own courses and for individual self-paced learners.

The decision was made to use content and resources from the Digital Library where appropriate in order to ensure eLGrid can:

- access materials created by others
- benefit from updates to the material by others
- quickly create new materials.

This ties in with calls for increased sharing of materials and stress on importance of digital libraries as laid out in, for example, [7].

The Digital Library stores metadata about resources in Dublin Core format [174] and stores information for each learning resource about the subject, author, provenance, format, suitable audience, etc. The library front-end implements search functionality to search for relevant materials on a range of metadata fields.

The course construction tool used in eLGrid, ACCT, provides support for automatic searching of a learning resource repository in an XML Database. While the ICEAGE Digital Library does not publish metadata in a compatible format, it seemed likely that a service could be created that would convert the metadata in the ICEAGE Digital Library to an XMLDB format in order to allow automatic searching for suitable learning content via the course creation tool. The majority of materials in the ICEAGE Digital Library are not, however, sufficiently fine-grained for use in the eLGrid system without some modification. These materials are largely in the form of Powerpoint presentations which span many 'concepts' within an adaptive eLearning course. For eLGrid it is necessary to have more fine-grained content, which can support a one-to-one or one-to-few mapping between course concepts and content.

Instead, suitable resources were found manually using the ICEAGE Digital Library front-end's search functionality, then modified and republished on a local website in the form of finer-grained html pagelets. The associated metadata was modified using a simple script created by the author which generated an XML file to be loaded into an XMLDB in order for it to be made available to APeLS and ACCT.

4.11.2 People Registry

ICEAGE also produced a Grid People Registry [95] which contains a list of educators and students. This tool has been used by several projects including EGEE and the International Schools of Grid Computing. The EGI project is now investigating the use of the People Registry to store information about trainers within each National Grid Infrastructure.

The Grid People Registry stores metadata about interests and areas of expertise of educators. Students who participated in (or even applied to) events run by ICEAGE, such as ISSGC/IWSGC, have a profile created in this system and this stores metadata about the modules they have completed as part of the school, becoming a type of e-Portfolio [175].

The possibility of standardising and sharing metadata between the Grid People Repository and the eLGrid learner profile was investigated, but issues of privacy and data protection prevented a solution being implemented. This might be revisited in future if a procedure could be devised for obtaining permission to share information about people who create profiles in the Grid People Registry.

4.11.3 Networking and Grid Services

It may not be feasible or desirable to replicate all of the Grid and network services of a production infrastructure within the t-Infrastructure. Where services are not to be replicated, external production services can be utilised, with the appropriate networking modifications to allow access to these services. In the prototype implementation for this thesis services such as NTP are not replicated, and the University's production NTP service is used instead. The networking layer is configured to give access only to these essential services. For more detail see Appendix B.

4.12 Courses

For the prototype implementation of this thesis eight courses were created using the tools described in Appendix C. Seven were created by the author and one, the *SQL* course, is a republished version of a course created by the Knowledge and Data Engineering Research Group in Trinity College. Figure 4-3 shows the login screen with a list of available courses. Each of these are discussed below and a summary is provided in Table 4.3

4.12.1 *Grid Certificates* course

The first course to be created was the *Grid Certificates* course which takes a new Grid user through the process of applying for and obtaining their Grid credentials.

The course provides a gentle introduction to the topics of authentication and authorisation, encryption and digital certificates as well as some of the commonly used Grid tools for authentication and authorisation, such as VOMS [176] and MyProxy [177].

The practical section guides the learner through the process of applying for, obtaining and using their Grid credentials. The interface to the eLGrid CA which the learners use to apply for a certificate is embedded within the eLGrid eLearning ap-

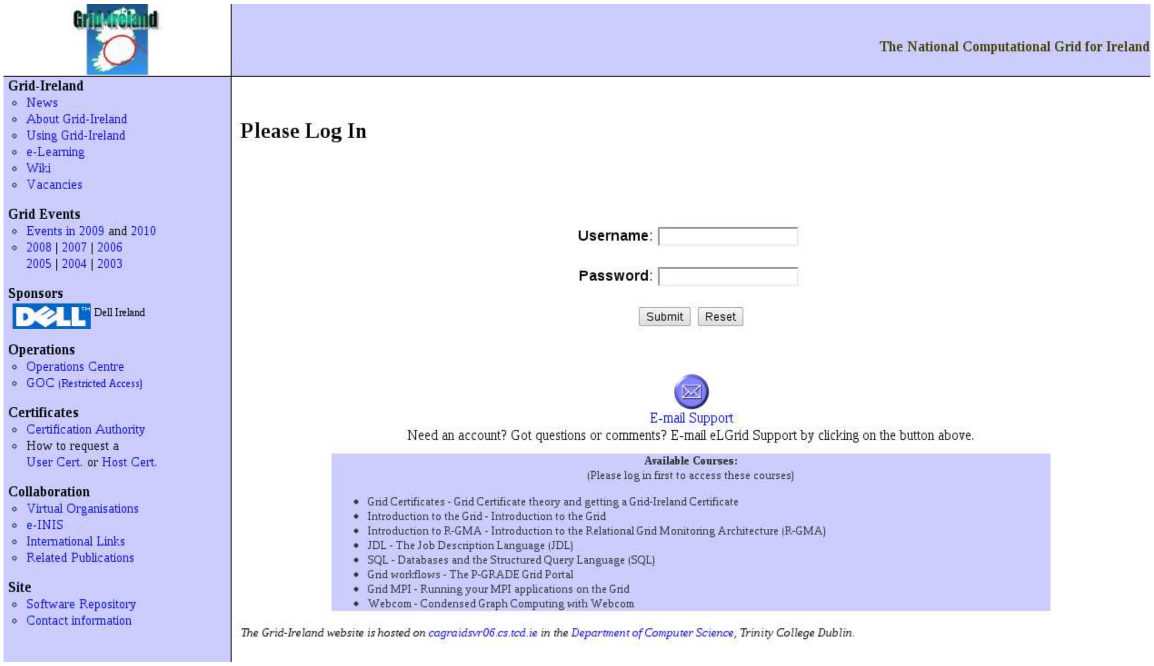

Figure 4-3: The eLGrid Login screen showing the list of courses

plication as shown in Figure 4-4, and is virtually identical to that of the Grid-Ireland production CA. The course integrates either the eLGrid CA or the Grid-Ireland production CA as appropriate and the learner can be adaptively presented with the correct environment to request their certificate.

The course contains 33 theory and practical-based concepts as well as a range of external resources providing more information, guidelines on selecting a strong password, links to the Grid-Ireland production CA, etc. 45 learners have begun the Grid Certificates course with over half having completed the practical section.

4.12.2 *Introduction to the Grid* course

The *Introduction to the Grid* course is loosely based on the curriculum of the International Summer School of Grid Computing (ISSGC). It aims to introduce the underlying motivations and challenges for Grid Computing as well as describing the typical components of a Grid. Some of the learning content is based on Powerpoint presentations given at the ISSGC, and some coarser-grained content is directly linked from the ICEAGE Digital Library, for example, video and SMIL [178] presentations

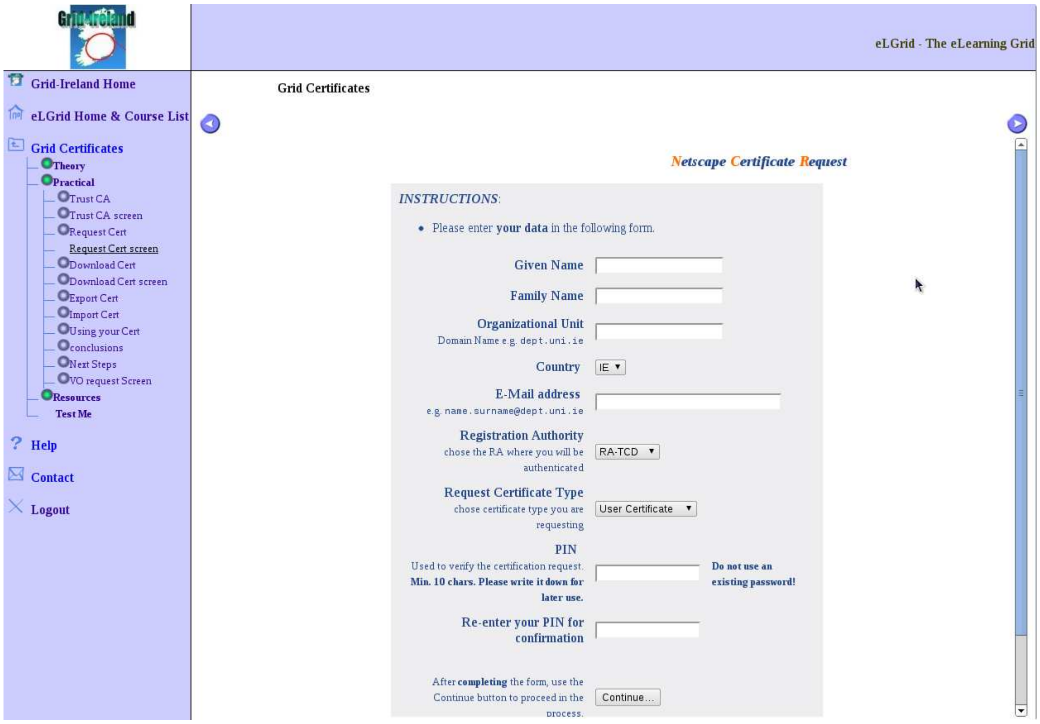

Figure 4-4: The eLGrid Grid Certificates Course practical screen

of introductory talks and interesting keynote speeches which took place at the school.

While some sections give middleware-specific examples, the course is not aimed at users of any particular Grid middleware, and the middleware-specific sections can be adaptively presented depending on information in the learner profile.

The course contains 48 concepts and has been attempted by 35 learners. Approximately one third have visited all of the core concepts and completed the necessary tests.

4.12.3 *Introduction to R-GMA* course

The *Relational Grid Monitoring Architecture* (R-GMA) [179] is a relational implementation of the Grid Monitoring Architecture defined by OGF [180]. R-GMA was developed within the DataGrid [181] and EGEE projects, but has been decommissioned for the European Grid Infrastructure [15].

The course contains 47 concepts, mostly theoretical. While it has been attempted by 19 learners, only three of those have completed the course with the others having viewed only a very small number of concepts.

4.12.4 *JDL* course

The Job Description Language (JDL) [182] is used by the gLite Grid middleware [183] to describe the requirements and properties of a job for submission to a gLite-based Grid. This course introduces the most commonly used features of the language, and enables the learner to create JDL files to describe single jobs and parallel parameter sweep jobs, with a focus on how to write simple scripts to manage the jobs.

The course contains 36 concepts describing the JDL language and syntax. 21 learners have attempted the course with approximately one third having completed it.

4.12.5 *SQL* course

This course introduces databases and the Structured Query Language (SQL) [184]. It is based on an existing SQL course produced by the Knowledge and Data Engineering Group in Trinity College [64] and reuses the concept space, narrative model and content from that course. The SQL course contains 114 concepts in total. 14 learners have attempted the course, but the completion level is low with only five learners having viewed the necessary concepts.

4.12.6 *Grid workflows* course

The P-Grade Grid Portal is a web-based graphical tool to create and run workflow jobs on the Grid. The course consists mostly of practical exercises with each of the relevant P-Grade screens embedded into the eLGrid eLearning front-end to allow the learner to access the practical environment from within the eLearning environment.

The course contains 28 concepts including the P-Grade screens. 36 learners have attempted the course with approximately two thirds having completed it.

4.12.7 *Grid MPI* course

MPI stands for Message Passing Interface [185] and refers to a standard API for passing messages between parallel processes. It is used widely for tightly-coupled parallel applications which need to interact, to synchronise their actions or to share information as they run. The course is specific to running MPI applications on the Grid, and assumes familiarity with MPI.

It contains practical exercises using four different Grid user-interface tools: the command-line UI tools, the Ganga UI, the Migrating Desktop, the Ganga graphical user interface, and the P-Grade portal.

The course contains 46 concepts, and has been attempted by 36 learners about half of whom have completed the course.

4.12.8 *WebCom* course

WebCom is a system to execute programs according to the Condensed Graphs (CG) model of computing [186]. This is an implicitly parallel graph-based model, similar to the dataflow computing model, but supporting availability-driven, demand-driven and control-driven execution strategies. WebCom-G is a Grid-enabled implementation of WebCom which submits users condensed graph jobs to the Grid [187].

This course introduces the theory behind the Condensed Graph model of computing before describing the WebCom interfaces and how to use them, along with a number of practical exercises. Learners must download the eLGrid node, a special node which must be included in their graphs in order to allow the results of their jobs to be automatically sent to the eLGrid eLearning application.

The course contains 63 concepts and has been attempted by 42 learners, all but three of whom have completed the course. The high completion level for the WebCom course is likely due to its use as part of an undergraduate module on Advanced Computer Architectures run as part of the bachelors degree in Computer Science at Trinity College Dublin.

4.13 Discussion

This chapter has described the XeL architecture and the prototype implementation, eLGrid, developed for this thesis. In doing so it has successfully answered a number of questions of this thesis:

1. *Is this instrument feasible?* Yes, the feasibility of such an instrument has been both demonstrated and described. The feasibility of the eLGrid t-Infrastructure has been described in section 4.3, while the feasibility of the adaptive eLearning application has been described in section 4.5.

2. *Can one provide distributed pedagogically-sound self-paced adaptive Grid education via eLearning?* Yes, an adaptive eLearning application has been developed along with a body of courses teaching Grid topics. Considerable effort has gone into the design and the choice of tools to ensure that it follows principles of constructivism and active learning, and that the courses are pedagogically sound by using tried-and-tested course creation tools (see Section C.4) and by repurposing, where possible, existing materials created by experts in the field as described in Section 4.11.1. The potential for the integration of the t-Infrastructure and eLearning application to enhance the potential for active learning has also been explored, and several techniques have been implemented, including providing multiple, distributed practice sessions, providing opportunities to review progress through the course (including the practical aspects) using the annotated traffic-light menu system, and providing timely feedback on practical exercises.

3. *Could the same instrument be used both for remote self-paced learners and in face-to-face courses?* Yes, the instrument is suitable for use by remote self-paced learners, but it can, and has, also been used in several face-to-face courses, as Chapters 5 and 6 will show.

4. *Can a t-Infrastructure for practical exercises be integrated with an eLearning environment?* Yes, see sections 4.3 and 4.5. The t-Infrastructure can use a

fair-replica of a production Grid infrastructure and dedicated virtualised resources to guarantee quality of service with minimal hardware requirements. The isolation of the t-Infrastructure ensures security, while the the use of tried-and-trusted fabric management and system administration tools contributes to ease of use and configuration. The integration encompasses the idea of placing the t-Infrastructure practical environment close to the practical instructions in order to reduce cognitive load caused by changing between the environments.

5. *Is it possible for such an integrated eLearning tool to capture relevant information about practical exercises?* Yes, the integration of the eLearning application and t-Infrastructure allows certain information about the practical exercises and their environment to be captured. Modules have been developed which feed this information back to the eLGrid eLearning application as described in Section 4.7. The information can then be passed to the personalisation engine.

6. *Can the single sign-on capabilities of the Grid be leveraged for the eLearning application?* Yes, this requirement in section 4.4.1 was met via a certificate authority (see section A.1.2) and associated authentication mechanisms (see section 4.5.2). However, as Grid-Ireland does not yet support federated identity, neither does eLGrid.

7. *Can the learner's conduct of practical exercises be used by the eLearning personalisation engine?* Yes, the information about the learner's practical exercises captured by the XeL modules (see Section 4.7) is stored by the eLearning system and used to update the learner's profile with the potential to support a wide range of personalisation techniques, such as modification of the annotated traffic-light menu.

8. *Is it possible to treat satisfactory completion of practical exercises as pre-requisites and for assessment purposes?* Yes, the information about completion of practical exercises stored in the learner's profile is used in the same way as the *Test Me!* feature, allowing assessment of the learner's understanding of the

related concepts. Satisfactory completion can thus be used as a prerequisite to subsequent concepts.

9. *Can learners revisit adaptive eLearning including integrated practical exercises?* Yes, the eLGrid application facilitates the revisiting of concepts, including practical exercises by learners. Updates to the practical environments and tools can be reflected in the learner's personalised annotated traffic-light menu by the simple expedient of the tutor updating the learner's profile. Learners will thus be made aware that they should attempt the practical exercise once more. This could be extended in future to include email alerts, etc.

The requirements set out for both the t-Infrastructure (Section 4.4.1) and for the eLearning application (Section 4.5) have largely been met. Table 4.1 page 83 summarises how the t-Infrastructure design requirements have been met by the architectural design and the specific technologies chosen. Table 4.2 page 84 does the same for the eLearning application design requirements.

Some areas for improvement remain. The t-Infrastructure was created using the best technologies at the time of implementation, however it could benefit from the use of cloud technologies. Projects such as StratusLab [103] and DemoGrid [140] provide interesting technologies which might enhance the eLGrid t-Infrastructure. The use of results of the learner's practical exercises for personalisation could be greatly expanded beyond what is currently in place. Federated identity has yet to be explored.

Eight courses have been developed, as described in section 4.12. Table 4.3 summarises these courses and the XeL features which are implemented in each. Several courses display some level of integration of the eLearning application with the t-Infrastructure. For some this consists of a browser frame containing the web-based Grid user interface (web SSH client, web-based Grid portal, etc.) that can be displayed as part of a page in the eLearning application. Other courses, such as the *WebCom* course, include more advanced integration of the practical exercises, enabling relevant information about the learner's conduct of these exercises to be captured and used by the adaptive eLearning application.

Table 4.1: How the architecture and technologies used have met the requirements

Requirement	**Architectural Features**	**Technologies**
Integrate eLearning and Grid environment	XeL	APeLS, XeL modules, inclusion of practical environments within course navigation
Simulate a production infrastructure	Virtualised fair replica of production infrastructure	Xen, Network Configuration, Quattor, TransDeploy, GridBuilder
Enable single sign-on	Sign-on to the eLGrid eLearning application via Grid X.509 certificates	Tomcat, GridSite
Dedicated resources to guarantee QoS	Dedicated eLGrid infrastructure	Xen, Network Configuration, Quattor, TransDeploy, GridBuilder
Isolate training testbed for security	Firewalled off from production	Xen, Network Configuration with Iptables & Arptables
Make easy to use and configure	Customised fabric management and deployment tools	GridBuilder, TransDeploy, Quattor, APeLS

The implementation described here affirms the feasibility of the XeL approach and the proposed architecture. A number of experiments were also carried out in order to confirm this feasibility, and more specifically to explore whether XeL provides any benefit to the learner. The next chapters describe those experiments.

Table 4.2: How the eLearning tools & technologies used have met the requirements

Requirement	**Features**	**Technologies**
Principles of constructivism	Course construction method and design of features of the eLGrid eLearning application	ACCT, eLGrid eLearning application
Active learning	Opportunities for practice, review and reflection on learning	Integrated t-Infrastructure, the eLGrid annotated traffic-light menu system, ability for learners to review the results of their practical exercises captured through XeL
Pedagogically sound courses	Course construction process and design of concept space, narrative, and choice of learning content	ACCT, integration with ICEAGE Digital library
Support Personalisation of courses	Personalised adaptive eLearning system and courses created to take advantage of this	eLGrid annotated traffic-light menu system, APeLS and ACCT
Minimise extraneous cognitive load	Worked examples, multimodal materials, close-placement of practical environment and instructions	ICEAGE Digital Library, integrated t-Infrastructure, use of standards to enable sharing and re-use of material
Take advantage of the *distributed practice effect*	Distribute practical exercises throughout the courses	Integrated t-Infrastructure
Provide timely feedback	Integrate the learners' practical exercises and provide feedback	Integrated t-Infrastructure and XeL
Ease-of-use	Design the eLearning application to be easy to use	Adherence to best-practice guidelines for web development, use of the eLGrid annotated traffic-light menu system to support guided course navigation

Table 4.3: Summary of features of the eLGrid Courses

Course	**Concepts**	**Core Concepts**	**Learners**	**Non-zero Learners**	**Complete**	**Mean concepts viewed**	**XeL features**
Grid Certificates	33	20	45	37	17	10.49	Integration with the eLGrid CA
Introduction to the Grid	48	25	35	32	10	16.75	Integration with the eLGrid t-Infrastructure via a web SSH client
Introduction to R-GMA	47	32	19	11	3	12.27	None implemented, could add SSH access to eLGrid for practical exercises using R-GMA command-line tools
JDL	36		21	17	7	7.294	Integration with the eLGrid t-Infrastructure via a web SSH client
SQL	114	?	14	0	7	18.14	None implemented as this was a republished non-Grid-specific course
Grid Workflows	28	20	36	32	18	15.69	Integration with the P-Grade portal
Grid MPI	46	15	36	29	12	13.17	Integration with the P-Grade portal
WebCom	63	40	42	41	34	51.27	Results of practical exercises fed back from WebCom to the eLGrid eLearning tool
Totals	**415**	**152**	**248**				

Chapter 5

Evaluation of Grid eLearning without XeL

5.1 Overview

The first version of the eLGrid infrastructure went live in August of 2007 with a major upgrade in April 2008. Thus far 121 users have registered and over half of these have completed at least one course.

Five evaluation studies on the eLGrid infrastructure have been conducted as controlled experiments to date, three without XeL and two with. These took place during face-to-face courses and workshops held in Trinity College Dublin. Participants were asked to complete one or more eLGrid eLearning courses while tutors monitored their progress, answered questions and resolved any issues which occurred. The five experiments are summarised in Table 5.1.

There was also a cohort of learners who used the system without taking part in any evaluation study.

This chapter is concerned with the first three evaluation studies, i.e. those without XeL. It also looks at those learners who did not participate in an evaluation study.

The analysis presented here looks at the general usage of eLGrid without XeL. The purpose of the analysis presented in this chapter was to determine whether there are patterns as to how learners use eLGrid. Where learners had completed an end-of-

Table 5.1: Experiments

Experiment	**Date**	**Num Participants**	**XeL-enabled?**
First evaluation study	May 2008	11	No
Second evaluation study	September 2008	7	No
Third evaluation study	February – March 2009	14	No
Fourth evaluation study	December 2009	10	Yes
Fifth evaluation study	December 2010	11	Yes

course survey the relationships between their satisfaction scores and usage patterns were also analysed. This analysis gives us a baseline with which to compare results of the experiments with XeL, presented in Chapter 6.

The eLGrid learners can be broken into groups based on whether they had attended a workshop or used the system on their own, and can also be divided to some extent by knowledge 'level' (undergraduate, postgraduate, academic staff, etc.), although this information is not available for all learners.

5.1.1 Methodology

The data gathering and analysis performed for this thesis is an exploration of what is possible as well as a method to obtain meaningful data for the evaluation of XeL. While a qualitative approach is more common in evaluating eLearning systems, and can give us some insights, a more quantitative approach was selected for this study. The quantitative approach to evaluating eLearning is more rare and thus valuable as an attempt to gather a corpus of data which, as well as furthering the goals of this thesis, may be of use to other researchers. The raw data and tables of the statistical results for all of the analyses presented here are given in Appendix L.

It is important to note that the samples used in these evaluation studies were not randomly selected, but rather, they were selected based on convenience. The first two experiments were organised with Grid-Ireland in conjunction with other organisations wishing to run introductory Grid courses for their users. These other organisations were mainly responsible for advertising the course and attracting and selecting participants. The third experiment (and subsequent experiments covered in chapter 6) were run as part of the undergraduate Computer Science course at Trinity College Dublin.

The experiments were realistic in that they involved actual use of the eLGrid application and t-Infrastructure as part of a real course. The experiment subjects were also a realistic sample of potential Grid learners and users; some were application scientists wanting to use the Grid while others were undergraduate students. No volunteers were solicited to use the system as part of an evaluation study, and all participants might be assumed to have had at least some motivation to complete the courses.

The experiments involved completing one or more courses using the eLGrid application, and then completing a post-course questionnaire which asked learners about their impressions of the course and the eLearning system. Learners were informed that their anonymised data might be used for evaluation of the course and the eLearning tools, and that this research was being conducted as part of a PhD thesis.

The sample sizes for each of the studies are quite small (see Table)

It is desirable that inferences might be made from the results found here, but the reliability of any such inferences must be assessed in light of the non-random nature of the sample selection procedure, and the relatively small sample sizes involved. However, the collection and analysis of this data is an important first step, and further experiments using larger sample sizes are proposed as possible future work in Section 7.2.

5.1.2 The first evaluation study

For the first study eleven participants, all with an astrophysics background, attended the practical sessions held as part of the PIC simulations workshop organised by the Dublin Institute for Advanced Studies in May 2008 [188].

The users attempted four eLGrid courses: Grid Certificates, Introduction to the Grid, Grid Workflows and Grid MPI.

A paper questionnaire was administered to gauge satisfaction with the system at the end of each course.

Some technical issues (mostly to do with scalability of the web service) encountered during the workshop make the results of this first evaluation somewhat unreliable, but they may still yield some useful information.

5.1.3 The second evaluation study

The EELA project [134] organised a week-long Grid School in Dublin in September 2008 [189]. Seven users attended and all completed the Grid Workflows eLGrid course, with a few also attempting the Grid MPI course.

For this evaluation it was determined that a more quantitative method of assessing participants' learning was required. An online pre-course questionnaire was introduced to identify the learner's experience with Grid technology. The paper-based post-course satisfaction questionnaire used in the first evaluation study was also administered after the course.

Learners completed a number of practical exercises as part of the workshop. The level of completion is used as an objective measure of learning achieved. Based on these results it was possible to conclude that learning had taken place during the course, although it was not possible to precisely quantify that learning.

The complete pre- and post-course surveys and responses, along with the results of the practical exercises are presented in [190].

5.1.4 The third evaluation study

The third evaluation study was conducted on a group of fourteen undergraduate students of Computer Science in Trinity College Dublin in February-March 2009 [90]. The students were asked to complete the WebCom [187] eLGrid course (see Section 4.12.8) in their practical sessions with a view to using WebCom in their practical assignment.

A pre-course questionnaire and paper-based post-course satisfaction questionnaire were administered. The pre-course questionnaire was online, as for previous evaluations. Students were also assigned marks for their group practical projects and for their usage of the eLGrid system (based on responses to the adaptive tests). These marks give us an objective measure of the learning achieved using the eLGrid system that is more quantitative than the measure used in the second study.

5.1.5 System Logs

Learner activity is logged automatically by the system, yielding various metrics such as number of logins, number of times course material is viewed, number of tests taken and passed, etc. There is also timing data showing how long learners have spent viewing course materials.

This data is available for all learners, whether or not they have taken part in one of the three evaluation studies mentioned above.

5.1.6 Survey answers

The majority of learners who took part in an evaluation study completed an end-of-course survey which gives information about their opinion of, and satisfaction with, the system.

This data is measured on a Likert scale [191]. Learners were asked to indicate how much they agreed or disagreed with a set of statements on a 5-point Likert scale as shown in Figure 5-1.

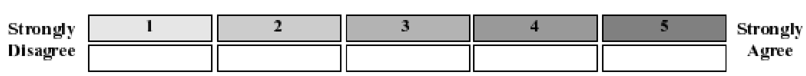

Figure 5-1: The Likert Scale for the post-course survey

The answers given on a Likert scale are inherently subjective and it is conceivable that a person who thoroughly understood or enjoyed an activity may give a mid-scale answer, while a person with a mediocre experience of the same activity may rate that experience higher [192]. Answering may also be subject to response biases such as *acquiescence bias*, where the respondent tends to give a neutral or positive response to all questions [193], or *social desirability bias* where the respondent respond in a way which will be perceived as being socially acceptable and/or desirable [194]. For example, the questions relating to whether learners found the course too hard could potentially suffer from social desirability bias if learners were unwilling to admit that they found the material difficult.

While not all learners have completed the survey (e.g. because they were not part of an evaluation study), there are also cases where a learner has completed the survey but left a particular question unanswered. This "No Answer" response is considered as a separate category of answer in our analysis and is not treated in the same way as the missing responses for learners who never completed the survey. While the Likert scale answers give us ordinal data which permits some limited numerical analysis, where a "No Answer" response is given it is a purely categorical data point and must be analysed as such. Thus "No Answer" responses are excluded from some of the more numerical analyses. For example, it is possible to calculate a statistic to describe the central tendency (median, mode, etc.) of the ordinal Likert scale survey responses, but it does not make sense to try to include "No Answer" responses in such an analysis. However, when combining Likert items to calculate an average satisfaction rating it may be more appropriate to exclude all data for learners with missing data points, or to use some missing value replacement technique [195].

Finally, while the "No Answer" category does not neatly fit into an ordinal view of the data, it might be plausible to interpret these responses as actually indicating disagreement with the statement presented. For example, if a learner leaves a blank for a statement such as "I understood how my test answers affected the traffic-light indicators" does this actually mean that he or she did not understand the question? If so is this evidence that the learner did not understand how their test answers affected the traffic-light icons? The problem of how to interpret "No Answer" responses is dealt with on a per-question basis and is documented in Appendix K.

5.2 Analysis

The R statistical package [196] was used for the majority of the statistical analysis and the statistical techniques used are all standard techniques available in the R package. The Weka data mining application [197] was used for the data mining.

The data for all learners was summarised and some general inferences made.

The learners were subsequently split into groups as shown in Table 5.2. To identify differences between these groups, t-test and one-way ANOVA F-tests were performed.

Table 5.2: Grouping of Learners

Group	Composition	Background
Group A	Participants in the first evaluation study	Physicists
Group B	Participants in the second evaluation study	Marine Sciences
Group C	Participants in the third evaluation study	Computer Science Undergraduate Students
Group D	Learners who did not participate in a particular experiment or study	Various

These tests assume a normal distribution of data, which does not always appear to be the case for the eLearning data, however, this may not be a huge problem for the analysis, since while the complete set of learners would not be expected to have normally distributed login counts (for example), the subset of learners who completed a particular course as part of a workshop might be expected to have a closer to normal distribution, with login count varying about the mean. The variation within each group would similarly be expected to be quite similar. The t-test and F-test are both relatively resistant to moderate non-normality and to unequal variances [198]. Thus these analyses have been applied in the hopes of learning something about the data, but the results must be interpreted with care. Furthermore, as the variances may not be equal across the groups the Welch t-test [199] has been used rather than the classic student t-test.

Where differences were found, multiple comparisons were used to identify which groups differed. Bonferroni adjustment techniques [200] were used to account for the effects of multiple comparisons on the probability of committing a type I error (finding a difference where none existed).

Groups A and B might be expected to show some similarities as both represent application scientists interested in using Grid technologies, although from different domains (physics and the marine sciences). Both groups also took part in short workshops in Grid technologies.

Group C consists of Computer Science Undergraduates using the system as part of their coursework and might be expected to show some differences from groups A and B, both because their use of the system took place over a longer period, and because their goals in using the system may have been different.

Group D are largely remote learners using the system over different periods and with different goals. One might expect to see the greatest variety within this group.

The learner's level (i.e. whether they are Undergraduates, Postgraduates, Post-doctoral researchers or Academic staff) is known for a subset of the learners. It is possible to perform some analysis of data classified by level, however because of the missing values the results are not always clear.

Where significant results are found in the analysis presented below, the results are reported in footnotes, giving the F and P values for ANOVA results and the alpha values for t-tests and multiple comparisons. Full results are tabulated in Appendix L.

5.3 Accessing the eLearning application and the learning materials

The information in the eLGrid logs can tell a lot about how learners access the eLGrid system. In particular, the number of logins, the number of times learners viewed course concept materials and the time spent viewing materials are all logged. These three measures taken together might give some idea of the usage patterns of the system.

(a) Number of logins: The maximum number of logins is actually quite low at 28 and the mean and median are both very low at 5.27 and 3 respectively. So most learners do not log into to the system very often. In fact, many users may have requested accounts more out of curiosity than anything else and may only have logged in once or twice to look around.

(b) Concept count: The concept count measures the number of times that learners viewed course concepts. The mean count is 84.96 with a standard deviation of 97.16. The median of 50.5 is lower than the mean, suggesting possible outliers, however, there is considerably less variation in the number of times concepts were viewed than in the time spent viewing.

(c) Mean time viewing concepts: The mean time spent viewing concepts is approximately three and a half hours. However, many of those learners who have only logged in a few times have not completed any courses and have viewed very few concepts. Conversely, a few learners spent a lot of time logged in, but were perhaps not always actively using the system, for example, the maximum time spent viewing concepts is over 80 hours, but as this particular learner remained logged in for long periods while only occasionally loading a new page, it seems likely that they may have been doing other things at the same time as reading the course materials.

5.3.1 Statistical Inferences

The three variables above were analysed by experiment group and by level, performing an ANOVA F-test to see if there were differences between the groupings.

Concept Count

In the case of concept count, the ANOVA F-test indicated that there was a difference between some of the groups[1]. Multiple comparisons reveals that group C has statistically significant differences from the three other groups. No differences between the other groups could be found[2]. Intuitively this makes sense as the first two groups took part in short workshops, while the third used the system over the course of several weeks and the fourth are made up of mainly remote learners. One might also have expected group D to show some differences from groups A and B, but no such differences were found.

[1] $F(3,70) = 12.183$ with $p < 0.001$

[2] $\alpha = 0.05$, $t_{A-C} = 4.049$, $t_{B-C}=2.944$, $t_{C-D}=6.049$

Similarly, when comparing the concept counts of the different levels, a statistically significant difference was found[3]. This analysis included only the subset for whom their level was known, but the results are similar even if the unknowns are included as a separate grouping. Multiple comparisons showed that the academic and postdoc concept counts both differed from the undergraduate concept count[4].

In this case there is no difference found between the undergraduates and the postgraduates. This is interesting as they are the closest in level and might reasonably be expected to exhibit similar behaviour. When they are compared in their experiment groups there is a difference between group C and the groups containing postgraduates (group A and D), however when the postgraduate learners are isolated and compared directly to the undergraduates this difference is not found. The difference is likely to arise from the other learners in those experiment groups and there is no reason to think that the postgraduates differ substantially from the undergraduates.

The undergraduate group also has the highest concept count. Not only does this group contain the highest individual value, but the median of this group is also the highest of all the groups, and in fact the minimum value is also above the medians of any of the other groups (see Figure 5-2). Most other groups included some learners with a total concept count of 0. In the case of the undergraduate group, however, participation in the eLearning course was a mandatory part of their undergraduate studies, which may explain the high minimum when compared to the other groups.

The concept count of the postgraduate group shows a relatively wide range and interquartile range (IQR), but has no outliers, while the academic and postdoc groups have quite a small range but some outliers. Of course, it must be remembered that the academic and postdoc groups are the smallest groups in the sample; if there were more data points in these groups one might see a different pattern, for example, the outliers might not actually be unusual values, and more data points might fall between the median and the outliers, thus changing the range and IQR substantially and in effect, transforming the outliers into normal data points.

[3] $F(3,30) = 7.363$ with $p < 0.001$

[4] $\alpha=0.01$ $t_{academic-undergrad}=3.861$, $t_{postdoc-undergrad}=3.717$

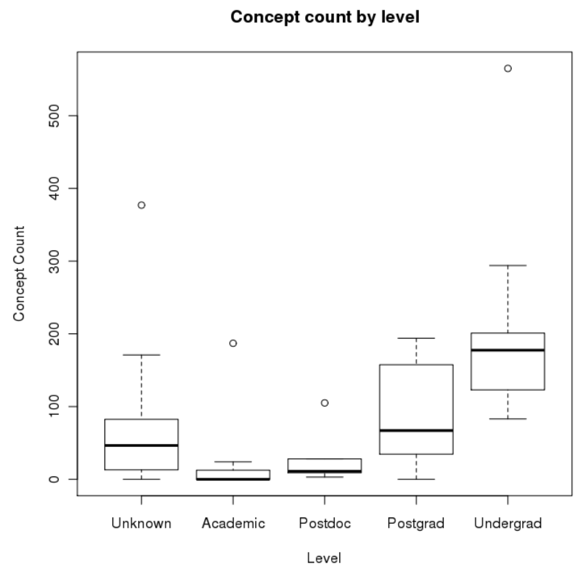

Figure 5-2: Concept count grouped by level

Login Count

When an F-test was performed on the login counts for the four groups, significant differences[5] were found, and multiple comparisons once again showed that it was group C which differed[6].

The same is true for login count viewed by level; a statistically significant difference was found between the undergraduates and all the other levels (this time including the postgraduates)[7].

The large variance and range illustrates the variety of learner behaviours encountered, but it is interesting to note that it may be possible to distinguish those who took part in short workshops from those who used the system over longer periods or from those who were remote learners.

Time viewing concepts

There was insufficient evidence to infer any statistically significant differences in time spent by the different populations when comparing learners by either experiment

[5] $F(3,70) = 14.262$ with $p < 0.001$

[6] $\alpha=0.01$, $t_{A-C}=3.560$, $t_{B-C}=5.237$, $t_{C-D}=5.932$

[7] $\alpha=0.05$, $t_{academic-undergrad}=5.068$, $t_{postdoc-undergrad}=5.57$, $t_{postgrad-undergrad}=3.227$

group or by level. Similarly no statistically significant difference was found in times when grouping the learners by level. There were, however, statistically significant differences in the ANOVA F-test results for the logs of time spent when broken out both by the experiment groups[8] and by the learner's level[9].

When multiple comparisons are done on the logs of time spent viewing concepts by experiment groups, groups B and D and groups C and D show differences. It suggests that there may be something different about the logs of time spent for group D[10].

For the logs of time spent viewing concepts by level there were differences between the academics and undergraduates, the postdocs and postgraduates and the postdocs and undergraduates[11]. Here again, the undergraduates are different from both the academics and postdoctoral staff. There is no evidence for a difference between undergraduates and postgraduates, but it can be seen that the postdocs differ from the postgraduates.

One might tentatively propose that academics and postdocs differ from postgrads and undergrads, while there is less evidence for a difference between the academics and postdocs, or between the postgrads and the undergrads. Postgraduates and Undergraduates may have more similarities to each other than they do to the academic and postdoctoral staff.

The undergraduate group by level contains the same 14 learners as experiment Group C. The other groups have a variety of levels and a number of unknowns. So while there is evidence for a difference between group C and all other groups for the number of concepts viewed and the logs of time spent viewing concepts when the data is analysed by experimental group, when looking at the data by level there is no evidence for a difference between this same group of learners (Group C or the undergraduates) and the set of postgraduates. This may suggest that postgraduates are more similar to undergraduates in their behaviour than either the postdoc or

[8] $F(3,60)=5.94$ with $p=0.001$

[9] $F(3,25)=10.389$ with $p < 0.001$

[10] $\alpha=0.05$, $t_{B-D}=2.853$, $t_{C-D}=3.692$

[11] $\alpha=0.05$, $t_{academic-undergrad}=3.367$, $t_{postdoc-postgrad}=3.828$, $t_{postdoc-undergrad}=4.793$

academic users.

However, to perform an analysis by level it was necessary to discard those learners whose level was not known, and so this analysis may not be as accurate as the analysis by experiment group, and the failure to find a difference between undergraduates and postgraduates may simply reflect this smaller sample size.

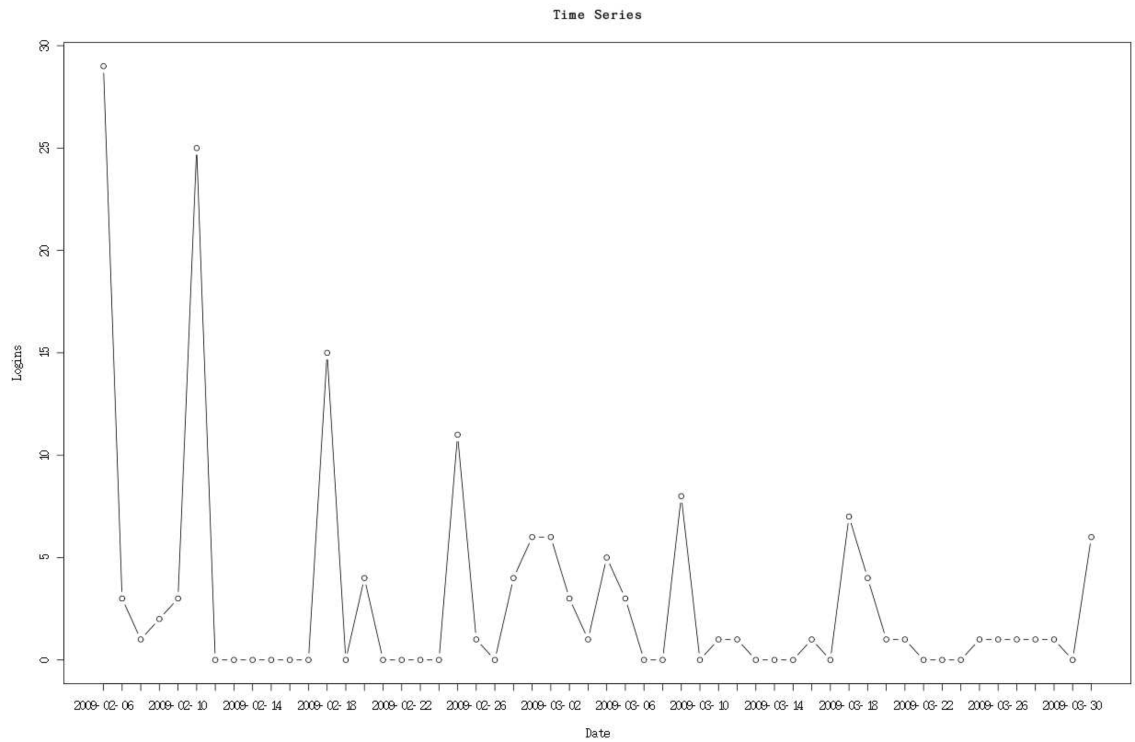

Figure 5-3: Time series plot of logins

Figure 5-3 shows this login pattern in the form of a time series plot. The first practical session was on the 6th of February 2009, and this corresponds to the largest peak as learners logged in for the first time. The next three peaks fall on 11th and 18th February, corresponding to the next two practical sessions. There are not very many logins between the scheduled practical sessions.

During the final week learners were working on their group projects and there was some activity on the system each day until the submission date of 6th March. The three smaller peaks after this date coincided with reminder emails sent to the learners asking them to be sure to complete any unanswered test questions.

It is possible to conclude that group C consistently spent longer using eLGrid, made more logins, viewed more concepts and took more tests than the other learners. The period over which they used the system was longer, however, as they had five supervised practical sessions of one hour each over five weeks. Many of the learners who attended workshops had rather longer practical sessions spread over the course of few days or a week, thus the overall practical time available to each group was similar. The differences seem to arise due to the undergraduate group's use of eLGrid outside of the practical sessions, which even though it was lower than that of use during the practical sessions, was still significant.

There is some evidence for differences by level, with those learners who are at a higher educational level (postdoctoral researchers and academics) spending less time and viewing fewer concepts than learners at lower levels (postgraduates and undergraduates) who appear more thorough in their use of eLGrid.

5.3.2 Correlations

A clear correlation might be expected between the number of concepts that a learner viewed and the time spent viewing concepts, and this was checked by plotting the data and performing a Pearson's correlation test. Concept count and time spent appear to be moderately correlated[12]. Interestingly, however, when looking at the correlations within the individual experiment groups there is a moderate to strong correlation for most groups but a very weak correlation[13] for group C (see Figure 5-4).

Strong and moderate correlations are also seen between the login count and concept count[14] and between login count and time spent viewing concepts [15]. Once again when looking at the individual groups it can be seen that group C exhibits moderate correlation between time spent viewing concepts and login count[16] while the other groups have a high correlation. It is not possible to tell at this stage whether there is any particular significance to this.

[12]Pearson's Correlation coefficient of 0.53

[13]Pearson's correlation coefficient = 0.16

[14]Pearson's correlation coefficient = 0.85

[15]Pearson's correlation coefficient = 0.64

[16]Pearson's correlation coefficient = 0.31

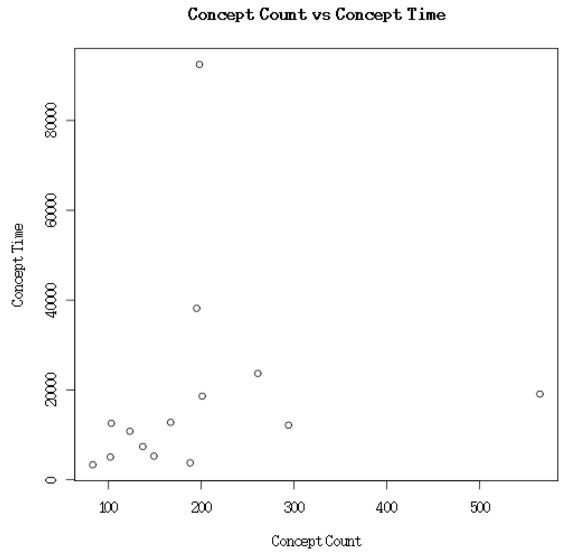

Figure 5-4: Concept count by time spent for Group C

It is possible to say that learners who returned to eLGrid regularly viewed more concepts on average than those who only logged in a few times, and those who spent more time on eLGrid also tended to view more concepts.

Table 5.6 page 120 summarises the main findings of the analysis of differences and correlations in learner usage patterns.

5.4 The *"Test Me!"* feature

The *"Test Me!"* link in the eLGrid course menu allows learners to take tests to check their progress. The tests not only give an indication of whether or not the learner has understood the material, they also feed into the adaptive "traffic-light" menu system. When a learner views a concept the traffic-light icon for that concept will change from green to amber. Once the learner completes the associated tests the icon will change to lilac.

eLGrid logs data about learners' use of the *"Test Me!"* feature. For example, the number of times the learner views the tests is logged along with the number of correct and incorrect test answers which they give.

The end of course survey also asks some questions specifically about the *"Test*

Me!" feature. Together these measure the pattern of usage of the tests and also the learner's opinions of and attitudes to the tests and their usefulness.

View Test Count is a measure of how many times learners clicked on the *"Test Me"* button in order to view the available tests. With a mean of 4.1 and a median of 1 it is obvious that most learners did not make much use of the tests. Some learners were much more active, and there is one particular outlier at 72 which is much higher than the majority of learners. Even excluding this learner there is still a fairly wide spread of values, as the next highest value is 29.

When an ANOVA F-test was performed, differences were found between the groups[17]. Further analysis using multiple comparisons shows that group C differs from the other three groups[18].

As can be seen from Figure 5-5, the undergraduate group (Group C) appear to have a much higher number of test views than any of the other groups, and this might be expected as they were required to complete the tests as part of their course.

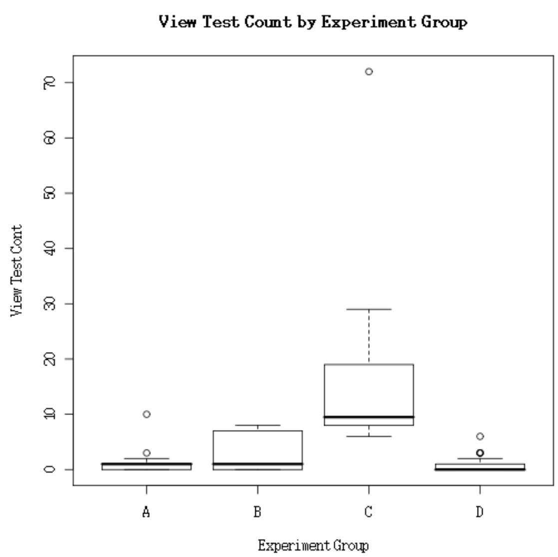

Figure 5-5: Test Count by Experiment Group

The *Correct Answer Count* measures how many times a learner answered the test questions for a particular concept correctly. In this case once again there are

[17] $F(3,70)=14.516$, $p < 0.001$

[18] $\alpha=0.01$, $t_{A-C}=5.187$, $t_{B-C}=3.636$, $t_{C-D}=6.406$

differences between the groups[19]. Multiple comparisons revealed that only group C was statistically different from the other three groups[20]. This is the undergraduate group, and again, this is not unexpected as this was the only group which had an incentive to complete the test questions.

There are also records kept of incorrect answers and incomplete answers (where multiple questions are required for a particular concept but the learner only answers one correctly). The ANOVA F-test did indicate that there were differences between the groups for incorrect answers[21] and incomplete answers[22], however, when multiple comparisons were performed these differences could not be identified.

When grouped by level, ANOVA again finds differences for both the number of test views and correct answers, but when multiple comparisons are performed only a difference between the view test count for academic and undergraduate groups is found.

The number of correct answers shows a moderate correlation with the total concept count (number of times learners viewed course concepts). A Pearson's correlation test produced a correlation coefficient of 0.57. This indicates that learners who view more concepts also tend to have more correct answers. There is also a strong correlation, however, between total concept count and view test count [23]. Thus it is not possible to say that learners who view more concepts necessarily do better at tests; it is more likely that they simply take more tests and thus have more correct answers.

Interestingly, while one might expect an inverse relationship between the number of correct and incorrect test answers, there is actually a strong positive correlation, as Figure 5-6 shows [24]. Those learners who had a high number of correct answers also tended to have a higher incorrect answer count. A likely explanation for this is that the learners with higher values for these variables simply took more tests than those with lower values; thus they had more opportunity to get both correct and incorrect

[19] $F(3,70)=12.222$ with $p < 0.001$

[20] $\alpha=0.01$, $t_{A-C}=4.741$, $t_{B-C}=3.983$, $t_{C-D}=5.755$

[21] $F(3,70)=12.687$ with $p < 0.001$

[22] $F(3,70)=14.688$ with $p < 0.001$

[23] Pearson's correlation coefficient = 0.74

[24] Pearson's correlation coefficient = 0.98

answers.

Plotting the same variables as ratios of the number of tests viewed does not show any particular pattern, and a correlation test finds no correlation[25].

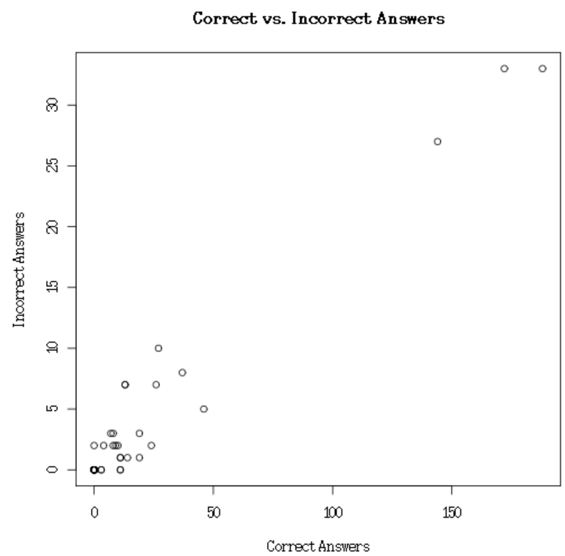

Figure 5-6: Plot of Correct versus Incorrect Answers

In the post-course survey there was one question which related specifically to the tests. Learners were asked to indicate their agreement with the statement "The tests were a good way for the system to evaluate my knowledge". Over 60% of learners agreed with this statement, with only 14% disagreeing. This suggests that learners generally felt the tests to be effective.

An ANOVA F-test showed that there is no statistically significant difference in response to this statement between the groups. Interestingly, responses were mainly positive whether the learners in the group had taken many tests or not. There was only a very weak correlation between the number of tests completed and the responses to this statement.

One might reasonably suppose that the values of variables such as *View Test Count*, *Correct Answer Count*, etc. might also be related to the survey responses to this statement, however, the data showed again only a very weak correlation between these variables.

[25] Pearson's correlation coefficient = 0.206

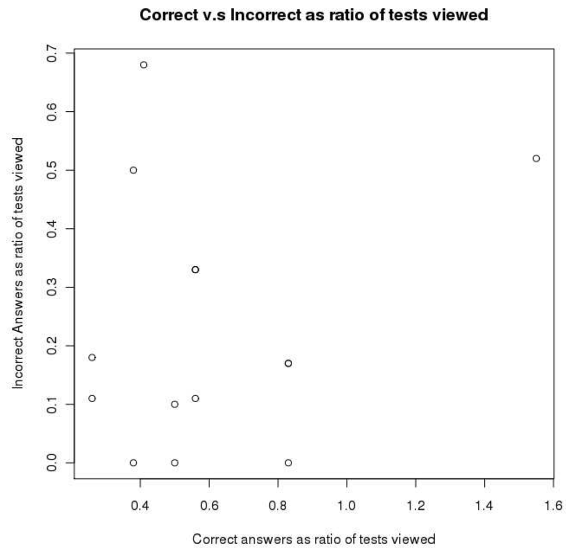

Figure 5-7: Plot of Correct versus Incorrect Answers as a ratio of tests viewed

5.5 Practical Exercises

The eLGrid infrastructure is used to teach users about Grid technologies. As such, many courses include a practical component where learners have an opportunity to try out their knowledge using the tools and environments about which they have been learning. This practical component is an important part of eLGrid which aims to integrate, as far as is possible, with the Grid environment in order to make it easy to complete the practical exercises.

The learner's attitudes about the practical exercises and environments are captured in the end of course survey. The statements relating to the practical exercises are shown in Table 5.3.

The responses were generally positive, with the majority of learners agreeing with all of the statements.

The mode (item with the strongest response) for each statement was *"4. Somewhat Agree"*, except for the statement "The theory in this course supported and complemented the practical exercises" where an equal number of learners choose *"3. Neither Agree nor Disagree"* and *"4. Somewhat Agree"* giving a mode of 3.5.

Items in a Likert scale can be combined to get an average satisfaction rating for

Table 5.3: End of course survey: Practicals

Number	Statement
1	The theory in this course supported and complemented the practical exercises
2	The practical exercises helped me to understand and apply the course concepts
3	The practical instructions were easy to understand
4	The practical environment was easy to access
5	The practical environment was easy to use
6	The practical exercises were a good way for me to evaluate my knowledge

the concept to which the items apply. In this case, the responses to all of these statements on satisfaction can be averaged to give us a mean satisfaction score for the practical exercises. The score calculated in this manner is 3.92, so on average, learners give a 3.92 satisfaction rating for the practicals on a scale from 1 to 5, with 5 being the most positive rating.

There is some controversy in the literature over whether analysis of variance techniques (ANOVA) can be applied to ordinal data such as these [201], however, one can choose to apply the F-test and merely be wary when interpreting the results. Adopting this approach, the majority of items showed no differences based on experiment group, and because not all learners had completed surveys and the learner's level was not known for all users, the sample size would have been too small to perform an analysis based on level.

ANOVA found statistically significant differences between groups only in the responses to the statement "The practical environment was easy to access"[26]. Multiple comparisons found that groups A and C were different, but there was not enough evidence to conclude that B and C were also different[27].

Figure 5-8 shows a box plot of this statement by group. Group A appears to agree more strongly than group C. This is unexpected because networking problems were encountered during the first experiment which caused some problems accessing one of the practical environments. The response may be more to do with the general

[26] $F(2,26)=3.632$ with $p=0.04$

[27] $\alpha=0.05$, $t_{A-C}=2.599$

attitudes and level of experience of the learners and their comfort-level with technology, as the physicists of whom group A was composed tend to be very technologically competent.

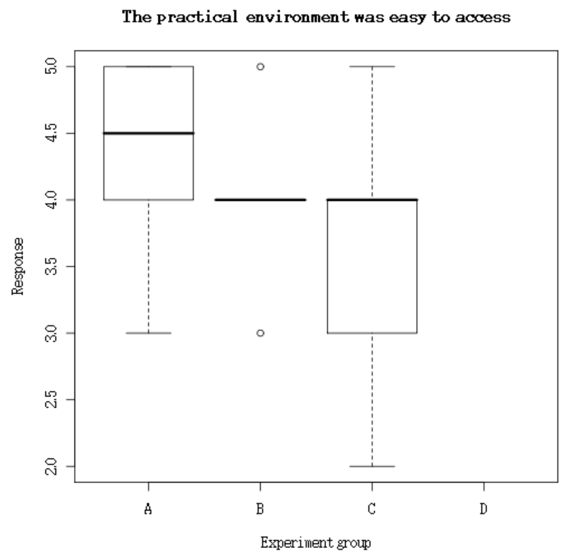

Figure 5-8: Plot of responses to the statement "The practical environment was easy to access" broken out by experiment group

Again, one should be wary of attaching too much importance to this result as the ANOVA techniques may not be particularly suitable for this variable, furthermore the responses are quite subjective.

These findings are summarised in Table 5.8 page 125.

5.6 Overall views of eLGrid

Some of the post-course survey questions give a general impression of how learners felt about the eLGrid eLearning infrastructure; "Was it easy to use?", "Was it easy to access?", etc. This information is very useful, highlighting areas where the infrastructure could use improvement.

Learners' attitudes to the adaptive traffic-light menu mechanism, which forms the core of the adaptivity and personalisation in eLGrid, were also surveyed, see the statements in Table 5.4. Positive attitudes were observed, which suggests that the learners do find the adaptation useful.

Table 5.4: End of course survey: Overall attitudes

Number	Statement
1	The course expected too much prior knowledge
2	The course was too simple, explaining things that I already knew
3	The presence of a tutor was necessary for me to complete this course, I would have had trouble if I was trying to use this course on my own
4	The eLGrid system is easy to use
5	The traffic-light indicators helped me to navigate the course
6	I understood how my test answers affected the traffic-light indicators
7	It would be useful if the traffic-light indicators took into account the results of my practical exercises as well as my test answers
8	In the future I will probably use the technologies taught in this course
9	If I need to learn about Grid technologies in future I would use the eLGrid eLearning system

In general learners appeared to be happy with the level of the courses with almost all indicating that the courses were neither too hard nor too easy.

Two learners indicated that their course was too easy, but as Figure 5-9 shows, these learners each had quite high values for both time spent viewing concepts and number of concepts viewed. One might conjecture that they found the course easy because of their thorough approach to reading the course materials while other learners who spent less time on the materials were less inclined to think the course was too easy, but of course, with only two learners in agreement with the statement it is not really possible to make any inferences.

The vast majority of learners agreed that the eLGrid infrastructure itself was easy to use, but a large number of learners nonetheless felt that a tutor was required in order for them to complete the courses. There is, however, quite a large spread in responses to this statement, suggesting that learners have widely different views on the subject. This response is not what was hoped for, as intuitively an eLearning system should ideally be capable of being used by remote learners without any tutor. There were some technical problems during one of the workshops where survey

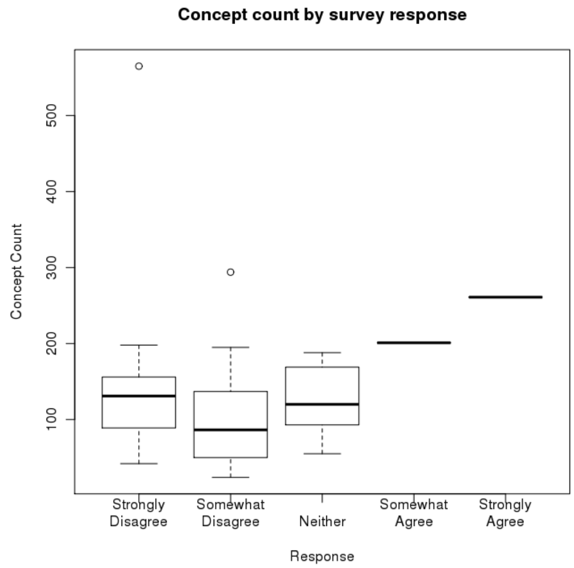

Figure 5-9: Plot of concept counts grouped by responses to the statement *"The course was too simple, explaining things that I already knew"*. Note that the two learners who agree with the statement have higher than average concept counts.

responses were gathered which might have contributed to these results. Nonetheless, eLGrid's content or user interface may need some further development in order to make it more suitable for remote learners.

In order to learn more about this result an ANOVA F-test was performed and statistically significant differences between the experiment groups were found[28]. Multiple comparisons showed that groups A and B differed from each other[29]. Group C appears to span the full range of responses.

Figure 5-10 suggests that Group A appears to be less inclined to believe that a tutor was necessary than group B, and this is also a surprising result. Networking issues during the first user evaluation meant that tutors were in fact require in order to resolve the issues and allow learners to progress through the course. The technical background and "can-do" attitude of the physicists of group A may be more of a factor in this response than the actual difficulty in accessing eLGrid. These learners may also have been discriminating enough to disregard the technical issues and judge whether

[28] $F(2,26)=4.688$, with $p=0.01$

[29] $\alpha=0.05$, $t_{A-B}=3.062$

a tutor was required simply in terms of the difficulty of the material presented.

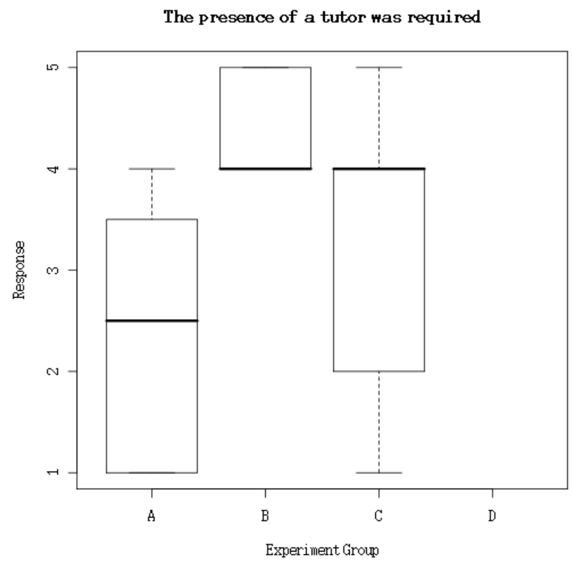

Figure 5-10: Responses to the statement “The presence of a tutor was required for me to complete the course”.

Responses to the statement “The traffic-light icons helped me to navigate the course” were quite varied, using the full scale from 1 to 5, however, the majority of learners agreed with the statement. The same is true for the statement “I understood how my test answers affected the traffic-light icons”. The mode for both of these statements was *“4. Somewhat Agree”* with a mean response of 3 or more. No differences between the experiment groups were found using ANOVA.

Interestingly there were some empty responses to these statements where learners answered the other questions but left these ones blank. It seems likely that learners who left these blank had not understood how the traffic-light indicators worked or how the test answers caused the coloured icons to change. These might thus be considered negative responses and would cause the overall response to be less positive.

When asked for their opinion of the statement “It would be useful if the traffic-light indicators took into account the results of my practical exercises as well as my test answers”, the majority of learners neither agreed nor disagreed. The incorporation of practical results into the adaptivity and personalisation of eLGrid was not yet available, but is considered in Chapter 6. The purpose of asking this question at that time was to ascertain the perceived need for this among users. These results indicate that the learners did not feel particularly strongly about the usefulness of such a feature. Even when analysing responses to this statement grouped by whether the learners found the traffic-light navigation based on their test answers useful, there was no evidence for a difference in responses between learners who liked or disliked the traffic-light navigation system. Thus in Chapter 6 it will be interesting to see if the perceived worth of such a feature differs from the perceived need once learners have actual experience of the system.

Finally, two questions asked learners to indicate whether they would use eLGrid or the Grid tools covered in the courses in future. There was general agreement that learners would use eLGrid again, but responses to the statement "In the future I will probably use the technologies taught in this course" was bimodal with equal numbers of learners choosing *"2. Somewhat Disagree"* and *"4. Somewhat Agree"*.

Only groups B and C had responses to this statement (the question was not asked in the first user evaluation study), thus a Welch two-sample t-test was sufficient to determine whether there was a difference between the two groups. The t-test confirms that the difference in responses of the two groups is statistically significant[30] and the size of the difference was found to be approximately between 0.1 and 2.7 points on the scale[31].

As always, when working with responses to survey questions, one must be cautious in interpretation. Because the scale is subjective and can differ from person to person, it should not be assumed that the points on the scale are all equidistant from one another. In this case, however, the difference seems large enough that it is safe to interpret it as a real difference in attitudes between the two groups, even if one cannot accurately interpret the confidence interval in order to determine how much of a difference exists.

While this t-test was two-sided, the plot shown in Figure 5-11 shows that Group B has a higher mean than Group C. This is as might be expected as group B contained scientists who were interested in using the Grid technologies presented, while Group

[30] α=0.05

[31] 95% confidence interval: [0.099,2.757]

C was comprised of undergraduate students only using the technologies because they had to as part of their course.

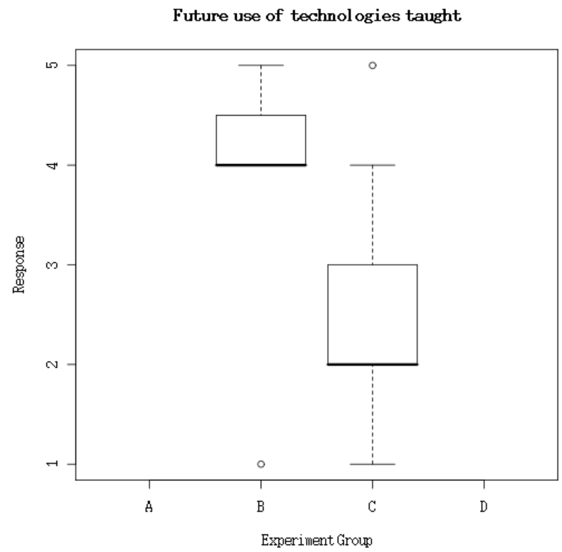

Figure 5-11: Responses to the statement "In the future I will probably use the technologies taught in this course".

Overall attitudes to eLGrid are summarised in Table 5.9 page 128.

5.7 Test & Project Marks: objective measure of learning

Three kinds of marks were assigned to learners in the third user evaluation group, i.e. those that attended the WebCom course. Table 5.5 explains the meaning of each.

The majority of learners correctly answered all of the eLGrid test questions, but some used trial-and-error to answer the questions for one difficult concept in particular. Learners who correctly answered all questions without using trial-and-error were given 100% and the majority got the full marks. The average eLGrid mark was thus very high at 95%.

The project mark was assigned based on adherence to the specification, completeness, appropriate use of the technologies taught in the course, and discretionary

Table 5.5: Test and Project Marks

Number	**Statement**
eLGrid mark	This was a mark assigned based on how many of the eLGrid test questions were correctly answered. The number of incorrect answers was also taken into account in order to ensure that learners who correctly answered the questions by trial-and-error were not rewarded.
Project mark	The mark assigned for the group project which learners completed.
Total mark	A weighted combination of the eLGrid and Project marks

marks for particularly good design or other factors. Most of the project groups produced relatively complete and working prototypes, and the marks were thus quite high, with an average of 75%.

A total mark was then calculated by combining the eLGrid and project scores, with more significance being placed on the practical project mark than on the eLGrid test mark. The result is an apparently normally distributed final mark with a mean of 83% and standard deviation of 8.22, as shown in Figure 5-12.

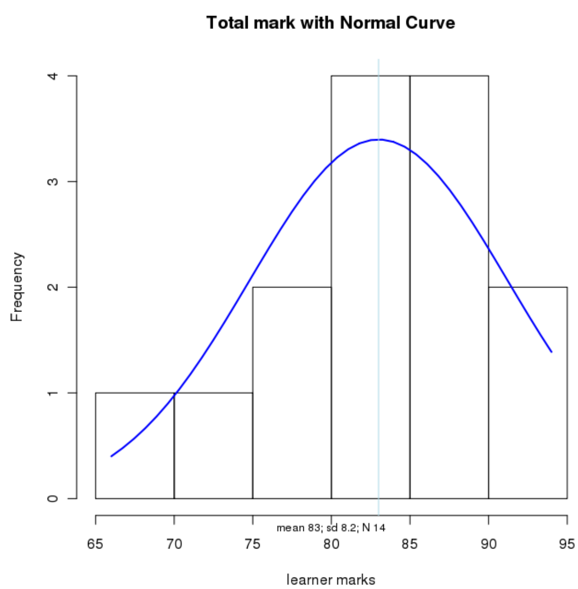

Figure 5-12: Distribution of learner's total marks, plotted with a normal curve.

As only the learners in group C have marks it is not possible to compare marks across experiment groups or levels. However, within group C one can look for correlations between high marks and particular patterns of behaviour. For example, a moderate degree of correlation is evident between the time the learner spent viewing concepts and their total mark [32]. Interestingly, however, a weak inverse correlation is seen between concept count and total mark [33], perhaps indicating that high concept counts (the need to study the concept many times) are correlated with weaker students. Stronger correlations were hard to find, however, and it was decided to apply some data mining techniques using the open-source Weka data-mining application [197] in order to see if any patterns could be found.

For example, using the method of *Attribute selection* [197] it was found that the most important predictors of total mark are the time spent viewing concepts, responses to the statements “The eLGrid system is easy to use”, “I understood how my test answers affected the traffic-light icons” and “The traffic light icons helped me to navigate the course”, and finally the project mark and eLGrid mark. The last two are used to compute the total mark so their importance is not surprising. A learner’s final mark appears also to depend on how long they use the infrastructure, how easy to use they find it, and their opinions and understanding of the traffic-light adaptive annotated menu system. This does not necessarily imply a causal relationship, and could simply be due to the better students finding other tasks, such as understanding and using the eLearning system, easier.

When the time spent viewing each concept was included in the analysis, it was found that three WebCom course concepts in particular are diagnostic of a learner’s final mark: the number of times the *“getArgDescription method”* concept is viewed and the time spent viewing the *“Running a graph”* and *“Operator execute method”* concepts. The *“getArgDescription method”* and *“Operator execute method”* concepts explain some of the methods which the learner must implement in order to create their own WebCom graph nodes, which was a required part of their practical project. The *“Running a graph”* node explains how to run WebCom Graphs and is obviously a key concept to learn in order to complete the practicals.

Test marks are only available for the final experiment and thus cannot be com-

[32] Pearson’s correlation coefficient = 0.35

[33] Pearson’s correlation coefficient = -0.37

pared across the experiments.

5.8 Data mining using course concepts

It might also be possible to find some useful information about usage patterns by performing data mining on a data set including the number of visits to, and time spent viewing, each concept of a course, rather than using the aggregate concept time and concept count variables which were analysed above. Each course, however, has its own unique set of concepts, so it is not possible to compare these details across courses. What is required is a set of learners who all completed the same course. Their data could be mined in order to identify clusters and possibly also decision trees and classification rules which might later be tested for their ability to correctly classify or predict behaviour for other learners using the same course. The identification of some of these rules has begun using group C and the WebCom Course, but further analysis using other cohorts of learners has yet to be done.

Using data mining software, a *One Rule classification* was performed on the responses to the statement “The tests were a good way for the system to evaluate my knowledge”.

The following rule was identified:

If time spent on the concept “Speculative Execution” $< 2349.5 \rightarrow$ agree
If time spent on the concept “Speculative Execution” $\geq 2349.5 \rightarrow$ disagree

This rule classified 12 out of 14 instances correctly.

Speculative Execution is quite a difficult concept, and was the one with which most learners in group C had trouble in the tests. Thus, the results show there might be a real relationship between the learner’s opinion of the tests and how successfully they learned this concept.

Both of the two learners who disagreed with the statement not only spent a long time viewing the Speculative Execution concept, but the majority of this time was spent towards the beginning of the course before they had attempted the test questions. Other learners tended to spend some time on this concept, attempt the questions, then return and spend some more time reading the material before correctly answering the test questions.

Clustering techniques were also used in the data mining tool to attempt to identify distinct clusters of learners. When asked to identify two clusters the tool split the learners into one group of 3 and another of 11. Comparing the learners in the two clusters it was apparent that while many of their scores seemed no different from each other, the learners in the cluster with 3 learners tended to have spent more time on the method definitions section of the course. This part of the course covered the Java methods that learners were required to implement in order to create their own WebCom nodes. Unfortunately there was no way to classify these learners by whether they had subsequently done well at the programming assignment. The fact that the project was a group project meant that the input of other learners outside of their cluster was relevant to their marks, and the learners did not achieve similar final scores. This classification process will perhaps be more useful when future learners are classified and their classifications compared to that of this group in order to see if the classifications are stable.

The data mining presented here has just scratched the surface of what might be possible. Further experiments would allow the results here to be compared and tested and a larger sample size might illuminate other interesting patterns in the data. This initial exploration of the application of data mining to the eLGrid data has given few clear patterns, but it has hinted at the potential worth of the techniques, and it suggests that there may be value in revisiting this in future.

5.9 Discussion

In spite of the small sample sizes involved, the analysis presented in this chapter has yielded several interesting results which suggest that there may be differences in the use of the system by different learner types. A number of tentative conclusions can be made about the eLGrid eLearning infrastructure based on this.

First of all, it is clear that there is a large variation in how learners use the system. Some learners spend a lot of time viewing concepts, some take a lot of tests, others return many times to view concepts, and so on. Some of the differences and correlations in how learners use eLGrid are summarised in Table 5.6 page 120.

There is no single pattern evident for how learners use eLGrid. This indicates that it is necessary to cater for different usage patterns. Indeed, because eLGrid is an adaptive eLearning system it should be able to handle the different demands of different learner types, but further investigation may be required in order to clearly identify clusters of learner types based on their usage patterns.

Interestingly there appears to be some evidence that learners at different knowledge levels use the system in different ways. There is clear evidence that group C differs in a range of behaviours from the other groups. Groups A, B and D contained a mixture of levels, and when compared to group C (the undergraduates) they exhibit differences across several variables. When learners are regrouped based on their level and a similar analysis is performed, differences can still be seen between the undergraduates and the academics and postdocs, but no evidence for a difference between undergraduates and postgraduates is found. This is potentially very interesting as it hints at a clear pattern in learners' use of the system with higher completion rates in most aspects at the lower levels, gradually decreasing as level increases. This might be used in course personalisation in future, by modifying the course content or navigation for learners at a higher knowledge level in an attempt to increase participation and completion rates. There is no evidence that learners from different scientific domains use the system differently.

It is important to note that there may be other factors contributing to these differences, such as steering by the tutors, or the group dynamic of the undergraduate group, which may have been different to the previous experiments. Further and more large-scale experiments would be required before definite conclusions can be drawn. The results do, however, suggest that it is worth exploring further.

In addition to these results, the analysis has proved fruitful for the evaluation of eLGrid. Face-to-face learners use eLGrid both within their practical sessions and outside as a self-paced learning tool. This suggests that eLGrid eLearning is suitable for use within face-to-face courses and that it facilitates additional learner self-study.

Attitudes towards the practical components of the courses, the *Test Me!* feature, and eLGrid itself were mainly very positive. 60% of learners agreed that the tests were a good way for eLGrid to evaluate their knowledge. There is little evidence of a correlation between the usage of the tests and attitudes to eLGrid, or that learners who take more or fewer tests fare any better in their final marks. Learners who spent more time overall using eLGrid were inclined to take more tests, and these learners tended to have high numbers of both correct and incorrect test responses, but it is clear that test usage is not particularly high, and learners will not take the tests unless they feel that there is some benefit to them.

Overall eLGrid seemed well-liked by the learners. The majority did not find the courses too difficult or too easy, and they largely felt that eLGrid itself was easy to use. In particular, the annotated traffic-light menu system appeared to be appreciated, with learners indicating that it helped them to navigate the courses. Learners also had a good understanding of how their test responses affected the traffic-light menu system. This is a positive indication for eLGrid and allows the analysis of the effects of XeL to proceed more confidently. However, as previously noted, one must be cautious in interpretation of these results as they are quite subjective and may suffer from various biases.

When asked for their response to the statement "It would be useful if the traffic-light indicators took into account the results of my practical exercises as well as my test answers", however, the majority of learners neither agreed nor disagreed. This suggests that learners do not see any great need for the annotated traffic-light menu system to take into account the results of their practical exercises (a core XeL concept). This unhelpful result (in terms of this thesis) will be compared with the attitudes towards the XeL version of the same course in Chapter 6.

The attitudes towards the practical exercises are summarised in Table 5.8 while overall attitudes to eLGrid are summarised in Table 5.9.

The analysis of the marks assigned to the students of course *4BA9* show that while learners who spend a lot of time viewing the course concepts tend to get higher marks, the inverse is true for the number of concepts viewed suggesting, perhaps, that weaker learners return repeatedly to the concepts while spending less time.

Further analysis using data mining revealed that the overall mark depended on how long learners use the infrastructure, how easy to use they find it and their opinions and understanding of the annotated traffic-light menu system. This suggests that the traffic-light menu system is, in fact, useful to learners and can help their learning.

Certain key concepts were identified using data mining, a similar approach might be taken with other courses to determine the key concepts for each and this information may be of use in course personalisation.

Data mining techniques also reveal that the time spent on certain concepts can be predictive of learners' responses to some of the survey questions. For example, learners who spent less time viewing the *Speculative Execution* concept are more likely to agree that the tests are a good way for eLGrid to evaluate their knowledge.

The *Speculative Execution* concept is an important but difficult concept in WebCom and a good understanding of this concept is vital in order to create recursive graphs, such as the graphs required for the learner's project. The majority of incorrect test answers by learners related to this concept. Learners who spent a lot of time here may have been having trouble with the concept and may thus have had many incorrect attempts to take the related test. However, further analysis of the learners who disagreed with this statement suggest that because they spent a lot of time on the concept at the start of the course before attempting the questions. Learners who agreed were more likely to have spent a little time on the concept then attempted the test before returning to review the concept when they had trouble with the test question. This difference in approach may affect learners' perceptions of the *Test Me!* feature. It would be interesting to explore areas such as this further in order to explore how learner style might affect their perceptions of online assessment.

Clustering techniques were able to separate learners into two distinct clusters, and again this might be further explored with future groups of learners to determine whether the classifications are stable.

Data-mining tools have the potential to identify many more learner patterns, but their real value is in creating classification rules in order to predict learner behaviour. This will require more experiments with learners using the same courses on which to test the rules which are produced.

Table 5.6: Differences in behaviour and usage patterns

Conjecture	Indication	Evidence	Confidence
Learners differ in their usage of eLGrid	Yes	large variance for most variables studied	High
Learner's knowledge level influences their eLGrid usage	Yes	Group C differs from all groups in Concept Count Academics and Postdocs differ from Undergrads in concept count Undergrads and Postgrads do not differ in concept count Academic and Postdocs do not differ in concept count Undergrads have highest mean concept count, followed by postgrads, postdocs and finally academics with the lowest mean value	Low
Remote learners differ from face-to-face learners	Perhaps	Log of time spent by Group D differs from all other groups, but no other differences could be found	Low
Face-to-face learners use eLGrid remotely to complement their class-based learning	Yes	Group C consistently show differences to other groups in several variables (e.g. concept count, time spent viewing concepts) Group C used eLGrid over a longer period, but overall contact-time was similar to other groups who attended short intense workshops Logins outside of class-time recorded, with peaks corresponding to submission deadlines and reminder emails	High
Learners from different scientific domains use eLGrid differently	No	No differences could be found between groups A and B but the sample size is too small to make reliable inferences	-

Continued on next page

Table 5.6 – continued from previous page

Conjecture	**Indication**	**Evidence**	**Confidence**
High concept count → high time spent viewing concepts	Yes	Most learners exhibit moderate to strong correlation Group C shows weak correlation	High
High login count → high concept count	Yes	Most learners exhibit strong correlation	High
High login count → high time spent viewing concepts	Yes	Most learners exhibit strong correlation Group C shows moderate correlation	High
More logins → more time and concepts	Yes	Correlations above	High

Table 5.7: Summary of the analysis of the *"Test Me!"* feature

Conjecture	Indication	Evidence	Confidence
Learners think that the multi-choice tests are worthwhile	Yes	60% agreed with the statement *The tests were a good way for the system to evaluate my knowledge* Only 14% disagreed	Moderate
Remote learners use the multi-choice tests	No	Remote learners had a mean view test count of 0.82 and a median of 0, suggesting that they do not make much use of the tests	High
General multi-choice test usage is high	No	Overall mean view test count of 4.1 and a median of 1	High
Learners only view the multi-choice tests if they have to take them	Yes	Group C who took the WebCom course as part of their college work and were required to complete the tests had the highest view test count variable, with a mean value of 16.29 and a median of 9.5. Group C also had the highest values for correct answer count and incorrect answer count	High
Learners did not take many multi-choice tests due to a low opinion of their value	No	The majority of learners responded positively to the statement *The tests were a good way for the system to evaluate my knowledge* whether or not they had taken many tests (60% agreed, only 14% disagreed) No correlation evident between view test count, correct answer count or incorrect answer count and responses to this statement.	High

Continued on next page

Table 5.7 – continued from previous page

Conjecture	Indication	Evidence	Confidence
Learners who view more concepts also view and take more multi-choice tests	Yes	Moderate correlation between total concept count and view test count [34], correct answer count [35], and incorrect answer count [36]	Moderate
Learners who spend more time viewing concepts also view and take more multi-choice tests	No	No correlation between total concept count and view test count, correct answer count or incorrect answer count	-
Learners who had many correct answers would have fewer incorrect answers and vice-versa	No	In fact there is a strong positive correlation between correct answer count and incorrect answer count suggesting that these learners simply took more tests	High
High concept count → high multi-choice test mark	No	Total total concept count and correct answer count moderately correlated, but total concept count and view test count also strongly correlated → Learners who view more concepts also inclined to take more tests and thus get more correct answers	-

Continued on next page

[34]Pearson's correlation coefficient = 0.74
[35]Pearson's correlation coefficient = 0.57
[36]Pearson's correlation coefficient = 0.55

Table 5.7 – continued from previous page

Conjecture	**Indication**	**Evidence**	**Confidence**
Learners who have incorrect answers take more multi-choice tests	No	Correct and incorrect test answers strongly correlated Learners who take more tests get more correct and incorrect answers → Learners did not simply retake the test because they had an incorrect answer, rather, these learners had a general tendency to take more tests	-
Learners who think the multi-choice tests are worthwhile take more tests	No	Learners responses to the statement "The tests were a good way for the system to evaluate my knowledge" generally positive whether they took many tests or not	-
Learners who get more correct multi-choice test answers are more likely to think the tests are worthwhile	No	No correlation was evident between the correct answer count or incorrect answer count and responses to the statement "The tests were a good way for the system to evaluate my knowledge"	-

Table 5.8: Summary of overall attitudes to practical exercises and environment

Statement	Results		
	Response	Frequency	Percentage
The theory in this course supported and complemented the practical exercises	Total	29	100%
	Strongly Disagree	0	0%
	Somewhat Disagree	1	3.45%
	Neither Agree nor Disagree	9	31.03%
	Somewhat Agree	9	31.03%
	Strongly Agree	8	27.59%
	No Answer Given	2	6.9%
	Mode: 3.5 equal numbers chose *"3. Neither Agree nor Disagree"* & *"4. Somewhat Agree"* Mean: 3.89 Median: 4 Range: 3		
	Response	Frequency	Percentage
The practical exercises helped me to understand and apply the course concepts	Total	29	100%
	Strongly Disagree	0	0%
	Somewhat Disagree	0	0%
	Neither Agree nor Disagree	5	17.24%
	Somewhat Agree	16	55.17%
	Strongly Agree	7	24.14%
	No Answer Given	1	3.45%
	Mode: *"4.Somewhat Agree"* Mean: 4.07 Median: 4 Range: 2		

Continued on next page

Table 5.8 – continued from previous page

Statement	Response	Frequency	Percentage
The practical instructions were easy to understand	Total	30	100%
	Strongly Disagree	0	0%
	Somewhat Disagree	1	3.34%
	Neither Agree nor Disagree	5	16.67%
	Somewhat Agree	13	43.33%
	Strongly Agree	10	33.33%
	No Answer Given	1	3.34%
	Mode: *"4. Somewhat Agree"* Mean: 4.1 Median: 4 Range: 3		
The practical environment was easy to access	Total	29	100%
	Strongly Disagree	0	0%
	Somewhat Disagree	3	10.34%
	Neither Agree nor Disagree	5	17.24%
	Somewhat Agree	15	51.72%
	Strongly Agree	6	20.69%
	No Answer Given	0	0%
	Mode: *"4. Somewhat Agree"* Mean: 3.83 Median: 4 Range: 3		

Continued on next page

Table 5.8 – continued from previous page

Statement	**Results**		
	Response	Frequency	Percentage
The practical environment was easy to use	Total	28	100%
	Strongly Disagree	1	3.57%
	Somewhat Disagree	3	10.71%
	Neither Agree nor Disagree	10	35.71%
	Somewhat Agree	10	35.71%
	Strongly Agree	4	14.29%
	No Answer Given	0	0%
	Mode: 3.5 equal numbers chose *"3. Neither Agree nor Disagree"* & *"4. Somewhat Agree"* Mean: 3.46 Median: 3.5 Range: 4		
	Response	Frequency	Percentage
The practical exercises were a good way for me to evaluate my knowledge	Total	28	100%
	Strongly Disagree	0	0%
	Somewhat Disagree	0	0%
	Neither Agree nor Disagree	5	17.86%
	Somewhat Agree	13	46.43%
	Strongly Agree	8	28.57%
	No Answer Given	2	7.14%
	Mode: *"4. Somewhat Agree"* Mean: 4.12 Median: 4 Range: 2		

Table 5.9: Summary of the overall attitudes to eLGrid

Statement	**Results**		
	Response	Frequency	Percentage
The course expected too much prior knowledge	Total	30	100%
	Strongly Disagree	15	50%
	Somewhat Disagree	8	26.67%
	Neither Agree nor Disagree	4	13.33%
	Somewhat Agree	1	3.33%
	Strongly Agree	1	3.33%
	No Answer Given	1	3.33%
	Mode: *"1. Strongly Disagree"* Mean: 1.79, Median: 1, Range: 4		
	Response	Frequency	Percentage
The course was too simple, explaining things that I already knew	Total	30	100%
	Strongly Disagree	10	33.33%
	Somewhat Disagree	10	33.33%
	Neither Agree nor Disagree	7	23.33%
	Somewhat Agree	1	3.33%
	Strongly Agree	1	3.33%
	No Answer Given	1	3.33%
	Mode: 1.5 equal numbers chose *"1. Strongly Disagree"* & *"2. Somewhat Disagree"* Mean: 2.07, Median: 2, Range: 4		

Continued on next page

Table 5.9 – continued from previous page

Statement	Response	Frequency	Percentage
The presence of a tutor was necessary for me to complete this course	Total	30	100%
	Strongly Disagree	6	20%
	Somewhat Disagree	3	10%
	Neither Agree nor Disagree	3	10%
	Somewhat Agree	11	36.67%
	Strongly Agree	6	20%
	No Answer Given	1	3.33%
	Mode: *"4. Somewhat Agree"* Mean: 3.28, Median: 4, Range: 4		
eLGrid is easy to use	Total	28	100%
	Strongly Disagree	0	0%
	Somewhat Disagree	2	7.14%
	Neither Agree nor Disagree	4	14.29%
	Somewhat Agree	15	53.57%
	Strongly Agree	7	25%
	No Answer Given	0	0%
	Mode: *"5. Strongly Agree"* Mean: 3.96, Median: 4, Range: 3		

Continued on next page

Table 5.9 – continued from previous page

Statement	Results		
	Response	Frequency	Percentage
The traffic-light indicators helped me to navigate the course	Total	28	100%
	Strongly Disagree	3	10.71%
	Somewhat Disagree	5	17.86%
	Neither Agree nor Disagree	5	17.86%
	Somewhat Agree	10	35.71%
	Strongly Agree	3	10.71%
	No Answer Given	2	7.14%
	Mode: *"4. Somewhat Agree"* Mean: 3.19, Median: 3.5, Range: 4		
	Response	Frequency	Percentage
I understood how my test answers affected the traffic-light indicators	Total	29	100%
	Strongly Disagree	2	6.9%
	Somewhat Disagree	5	17.24%
	Neither Agree nor Disagree	7	24.14%
	Somewhat Agree	10	34.48%
	Strongly Agree	2	6.9%
	No Answer Given	3	10.34%
	Mode: *"4. Somewhat Agree"* Mean: 3.19, Median: 3, Range: 4		

Continued on next page

Table 5.9 – continued from previous page

Statement	Results		
	Response	Frequency	Percentage
It would be useful if the traffic-light indicators took into account the results of my practical exercises as well as my test answers	Total	22	100%
	Strongly Disagree	1	4.55%
	Somewhat Disagree	5	22.73%
	Neither Agree nor Disagree	6	27.27%
	Somewhat Agree	4	18.18%
	Strongly Agree	4	18.18%
	No Answer Given	2	9.09%
	Mode: *"3. Neither Agree nor Disagree"* Mean: 3.25, Median: 3, Range: 4		
	Response	Frequency	Percentage
In the future I will probably use the technologies taught in this course	Total	21	100%
	Strongly Disagree	4	19.05%
	Somewhat Disagree	6	28.57%
	Neither Agree nor Disagree	2	9.52%
	Somewhat Agree	6	28.57%
	Strongly Agree	3	14.29%
	No Answer Given	0	0%
	Mode: Bimodal *"2. Somewhat Disagree"* & *"4. Somewhat Agree"* Mean: 2.90, Median: 3, Range: 4		

Continued on next page

Table 5.9 – continued from previous page

Statement	Results		
	Response	Frequency	Percentage
If I need to learn about Grid technologies in future I would use the eLGrid eLearning system	Total	21	100%
	Strongly Disagree	1	4.76%
	Somewhat Disagree	1	4.76%
	Neither Agree nor Disagree	3	14.29%
	Somewhat Agree	9	42.86%
	Strongly Agree	7	33.33%
	No Answer Given	0	0%
	Mode: *"4. Somewhat Agree"* Mean: 3.95, Median: 4, Range: 4		

Chapter 6

Evaluation of Grid eLearning with XeL

6.1 Overview

Two evaluation studies were carried out on eLGrid with XeL.

The same survey information was collected as for the evaluations without XeL, however the post-course survey statement which forecast the existence of XeL "It would be useful if the traffic-light indicators took into account the results of my practical exercises as well as my test answers" was removed, and the following statements that reflected the subsequent provision of XeL were added instead.

- "I understood how the results of my practical exercises affected the traffic-light indicators". The intent is to ascertain how clear the operation of XeL is to the learners.
- "The practical exercises were a good way for eLGrid to evaluate my knowledge". The intent is to identify learners' attitudes towards the effectiveness of XeL as a tool for adaptive assessment.
- "The fact that the traffic-light indicators reflected the results of my practical exercises was helpful". The intent is to see whether learners found XeL, and the

automatic updating of the annotated traffic-light menu system based on their practical exercises, a useful feature.

- "The instant feedback on the result of the practical exercises was helpful". The intent was to see whether learners appreciated the instant feedback on their practical exercises, and also whether XeL accelerated the learning process.

6.1.1 The XeL evaluation studies

The three evaluation studies described in Chapter 5 were carried out before the main XeL features were implemented (see Table 5.1).

Two further evaluation studies using XeL were carried out with the undergraduate students taking the fourth year course CS4021 *Advanced Computer Architectures* as part of the degree in Computer Science in Trinity College Dublin. This was a modified version of the course 4BA9 *Advanced Computer Architectures* in February-March 2009 which was used for the third evaluation study of eLGrid without XeL (see Section 5.1.4).

Unfortunately changes to the course (modularised from 4BA9 to CS4021 to reflect the Bologne Process [202]) led to less time scheduled for practical sessions, with only three one-hour sessions over the course of three weeks instead of the five sessions which had been scheduled for 4BA9. This meant that the scale of the project work undertaken by the students had to be reduced, and a more simplified project assigned. The content covered in the WebCom course was identical, and the aspects of using WebCom which the students had to demonstrate were also retained, however, the aim of the project work was simpler, involving the creation of a less complicated Condensed Graph workflow and implementing less functionality in the custom nodes that the students created. The group aspect of the 4BA9 project was also eliminated and each student completed the project on their own.

This change is unfortunate for the evaluation of XeL, as differences in the data between the groups may be attributed to the change in the structure of the course practical arrangements rather than the introduction of XeL.

The first XeL study took place in December 2009 and the sample group consisted of ten students. The second XeL study took place in December 2010 with a sample group of eleven students. None of the students repeated the course so the two samples contain different students.

The students were asked to complete the WebCom [187] eLGrid course in their practical sessions, including a number of practical exercises within the WebCom Course. Assessment again included multiple-choice questions in the *Test Me!* section of the course, but now also included automatic upload of results of the practical exercises. It is this XeL integration of the practical exercises with the courseware which is the main focus of the evaluation presented in this chapter.

A pre-course questionnaire and post-course satisfaction questionnaire were administered. The pre-course questionnaire was online, as for previous evaluations. The post-course questionnaire was paper-based for the first XeL study (as for previous evaluations), but was changed to an online format for the second XeL study. As with the third evaluation study without XeL (see 5.1.4) students were assigned marks for their usage of the eLGrid system (based on responses to the adaptive multiple-choice tests). The final practical exercise, which in contrast to 4BA9 was now delivered through eLGrid and included XeL, was also assigned a mark. An overall total mark was computed based on a combination of the eLGrid and practical marks.

The analysis presented in this chapter takes both of the CS4021 (experiments four and five) groups together as one cohort of XeL learners and compares them with the students who took the same Webcom course without XeL as part of the third experiment.

Some results are also presented from the first and second evaluation studies (see Sections 5.1.2 and 5.1.3 respectively) where they are relevant to the concepts of XeL explored in this thesis.

Once again, it is important to note that the samples used in these evaluation studies were not randomly selected, but rather were selected based on convenience. The sample consists of all students who took the course CS4021 *Advanced Computer Architectures* in December 2009 and December 2010. The sample sizes are also relatively small. It is desirable that inferences about the population of Grid learners and users might be made from the results found here, but the reliability of any such inferences must be assessed in light of the limitations of the experments.

6.2 Differences in learner behaviour with XeL

Although there was a difference in the length of the course module and the number of practical sessions between the 4BA9 and CS4021 courses, the patterns of usage of eLGrid are relatively similar.

Number of logins

The login counts for both groups are similar, with the XeL group having a slightly higher value. A t-test failed to find any significant difference between the login counts for the two groups, however, the XeL groups had fewer contact hours as part of their course than the non-XeL group, and a shorter period over which to complete the course, thus they might have been expected to have a lower login count and the difference may in fact be more significant than the figures suggest.

The outlier in the non-XeL group would fall inside the normal range of values for the group with XeL. This suggests that the small sample size of the non-XeL group may be a factor, a reminder to interpret the results with caution. The login counts for the two groups are plotted in Figure 6-1.

Login durations

The login durations are also similar as shown in Table 6.1. The mean and standard deviation show a large difference, but these are sensitive to outliers. A t-test found no significant difference between the means of the two groups. Median and inter-quartile range (IQR) are also shown in the table and these give a more accurate view of the average values and variation within each group. The XeL group has a median value just slightly lower than the non-XeL group.

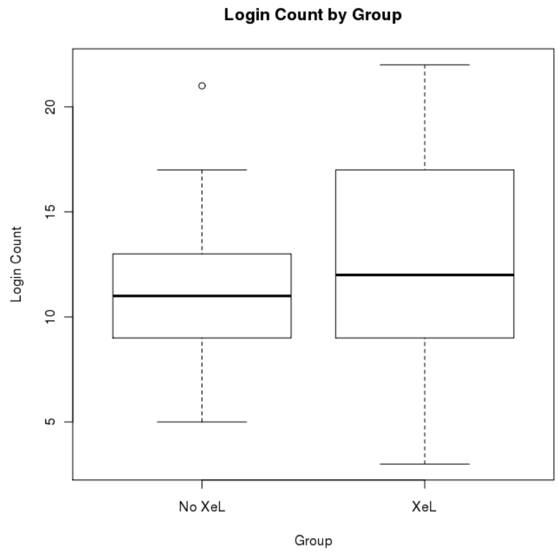

Figure 6-1: Login Count plotted with and without XeL.

Table 6.1: Summary of Login Durations for the non-XeL and XeL groups

Group	**Median**	**IQR**	**Mean**	**Standard Deviation**
No XeL	1210	4021	124400	1026420
XeL	1015	4511	11740	80187.79

Concept count and time spent viewing concepts

The total time spent viewing concepts and total concept count for the two groups are plotted in Figure 6-2 and Figure 6-3 respectively. These graphs show that the median value for the XeL group of both the total concept count and the time spent viewing concepts is higher than in the non-XeL group. The mean value for the total concept count is also higher in the XeL group, although the mean time spent viewing concepts is slightly lower (this is related to one outlier in the non-XeL group with a much higher than average time spent viewing concepts). The values are summarised in Table 6.2.

While the XeL values are mainly higher than those of the non-XeL group, the differences in these values is not very great and a Welch t-test failed to find any significant difference between the two groups (although again, the XeL group might

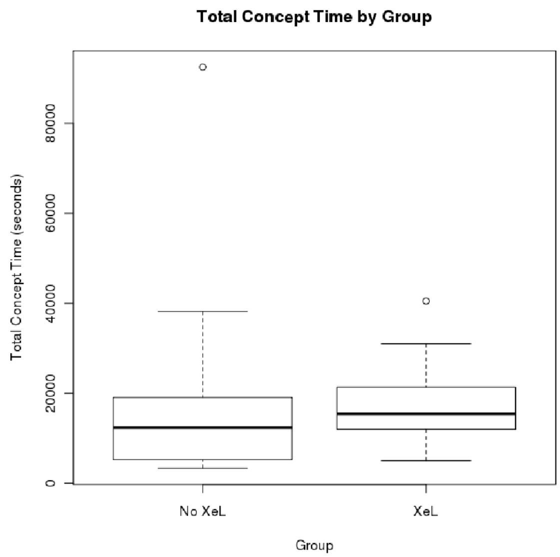

Figure 6-2: Total Concept Viewing Time plotted for the groups with and without XeL.

have been expected to have slightly lower values here).

How learners use eLGrid

There are, however, differences in the detail of how learners use the system. Where XeL-enabled practical exercises are available, learners seemed to repeatedly view the relevant course content (both the practical instructions and the related topics). Without XeL, learners appeared to gloss over many of the topics, and one might surmise that they ignored the practical exercises except where required to complete them as part of the course requirements.

This was true even where completion of a particular XeL-enabled practical exercise was not part of the final assessment for the course, and learners had been informed of this.

For example, the practical exercise on *Recursion* required learners to create a simple recursive condensed graph. The results were captured by XeL and stored in the eLGrid eLearning application, but did not contribute to the learner's *project mark* (see Section 6.3).

On counting the views of the Recursion practical exercise instructions and related

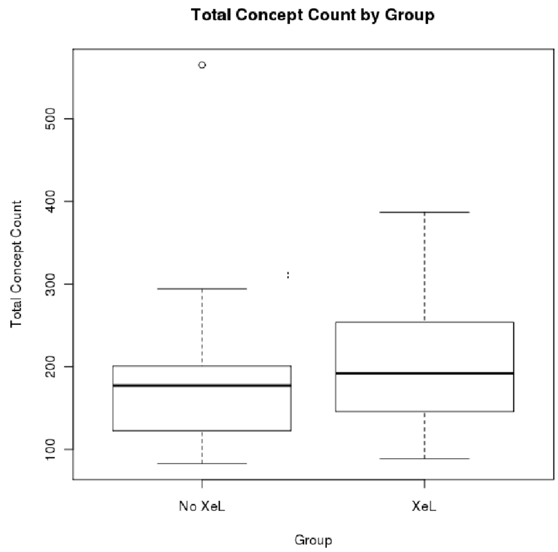

Figure 6-3: Total Concept Count plotted with XeL and without XeL.

concepts the group without XeL had a mean value of 5.21 viewings (standard deviation 4.53), while the XeL group had a mean of 14.36 (standard deviation 10.02). A t-test confirmed that there was a significant difference between the means of the two groups[1].

The pattern in learners' visits to these concepts is also interesting. A manual examination of the logs shows that the higher number of concept views appears to be due to learners revisiting the related topics after first viewing the practical exercise instructions, then returning to the instructions. In some cases this pattern was repeated more than once before the learner finally submitted their results by including the special eLGrid WebCom node in their graph. One can surmise that during this time learners are working on the practical exercise (without using the eLGrid node) and reviewing the materials.

Thus one can say that the presence of XeL appears to increase participation in practical exercises, while simultaneously encouraging review of related materials. This is likely to increase overall learning.

[1] α=0.05, 95% Confidence Interval: [2.105 12.514]

Table 6.2: Summary of Concept Count and time spent viewing concepts for non-XeL and XeL groups

Group	**Variable**	**Median**	**Mean**	**Standard Deviation**
No XeL	Total Concept Count	177.5	197.6	121.58
No XeL	Time spent viewing concepts	12350 seconds (approx 3 hours 25 minutes)	18920 seconds (approx 5 hours 15 minutes)	23178.74 seconds (approx 6 hours and 25 minutes)
XeL	Total Concept Count	192	205.5	73.1
XeL	Time spent viewing concepts	15420 seconds (approx 4 hours and 17 minutes)	17070 seconds (approx 4 hours and 45 minutes)	8237.127 seconds (approx 2 hours and 17 minutes)

Trial-and-error attempts

There is some evidence to suggest that the trial-and-error approach to the self-assessment multiple-choice tests may be reduced in the XeL group. While there is no evidence for a significant difference in the number of tests viewed or the number of correct answers, a t-test revealed a significant difference in the number of incorrect responses to the test questions[2]. Plots of the number of incorrect answers for each group are shown in Figure 6-4. The mean value of 1.33 for the XeL group is considerably lower than the mean of 5 for the non-XeL group.

Looking in more detail at the logs, it is possible to identify learners who used a trial-and-error approach based on the number of adjacent multiple-choice test attempts. If a learner made several attempts at a test with different answers, without reviewing the relevant concepts in between, it was considered to be a trial-and-error attempt.

Four out of the fourteen learners in the non-XeL group used a trial-and-error approach when answering some of the multiple-choice test questions, compared with

[2] α=0.05, 95% Confidence Interval: [-1.793 -15.398]

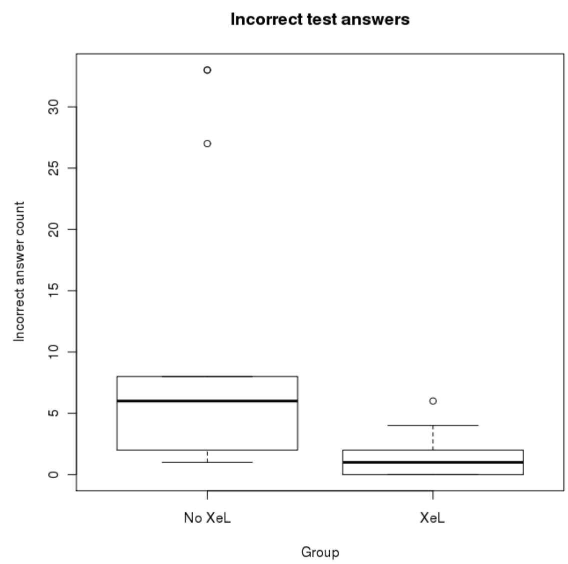

Figure 6-4: Incorrect Answer Count plotted with and without XeL.

only one of the 21 learners in the XeL group.

For the practical tests themselves, only one learner had incorrect responses, possibly suggesting that with XeL learners do not attempt to submit their practical exercise until they have fully tested it and are sure that it is working. Perhaps the presence of XeL encourages learners to take a similar approach to the multiple-choice test questions. This possibility is worthy of further investigation.

6.3 Evidence for increased learning with XeL

The work of the undergraduates who took courses 4BA9 and CS4021 was assessed and contributed to their overall mark for this course. A *project mark* (for the practical work), *eLGrid mark* (for the multiple-choice tests) and *total mark* were assigned. Table 5.5 describes the eLGrid and total marks. For the learners on course CS4021 using XeL, the project mark was based on their successful completion of the final eLGrid XeL practical exercises as well as the completeness and correctness of their code.

When a Welch two sample t-test was performed on each of these marks, a statistically significant difference was found between the means of the practical marks[3] and the total marks[4] of the two groups. The XeL group had significantly higher mean total and practical marks than the group without XeL. For the total mark, the difference was calculated as falling somewhere between 1% and 17%, while the practical mark had a difference calculated as between 16% and 29%. There was no significant difference between the eLGrid marks for the groups with and without the XeL features, thus it is fair to conjecture that any improvement in marks was due to improvements in the (XeL-assisted) practical marks.

These results appear to support the hypothesis that the application of XeL technologies to Grid eLearning increases learning. It must be remembered, however, that because of the change in the course structure between the (non-XeL) 4BA9 class and the (XeL) CS4021 classes, the practical aspect of the course was simplified and at least part of the improvement is likely due to this simplification.

6.4 Evidence for not needing a tutor with XeL

Comparing the responses to the statement "The presence of a tutor was necessary for me to complete this course" there is evidence for a significant difference between the non-XeL and XeL groups[5], with the XeL group more inclined to disagree. The responses to this statement for the two groups are shown in Figure 6-5.

In fact this is in contrast to all of the previous courses run using eLGrid in which the majority of learners consistently indicated that the presence of a tutor was required. This suggests that eLGrid with XeL is better able to cater for learners working without a tutor, or at a minimum, learners may believe that it is.

One might speculate as to the reasons for this. The instant feedback provided by XeL to learners may help them in completing their practical exercises, but as most learners only submitted their final results via the special WebCom eLGrid node once they had already tested their code and were confident that it was working, it seems

[3] α=0.01, 95% confidence interval: [16.373 29.342]

[4] α=0.05, 95% confidence interval: [0.987 17.108]

[5] α=0.05, 95% Confidence Interval: [0.079 2.195]

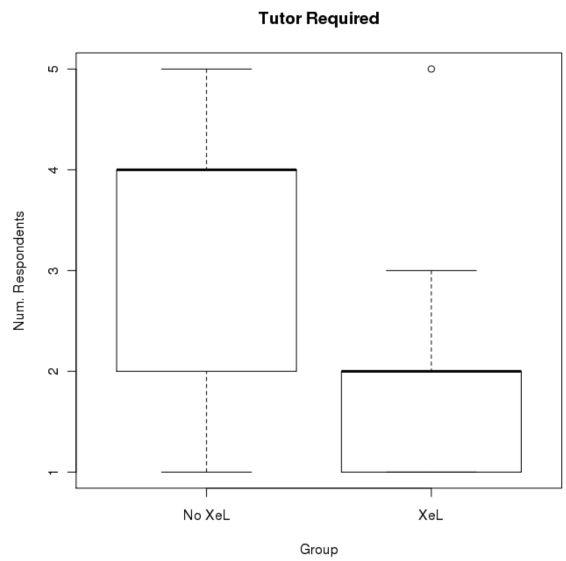

Figure 6-5: Responses to the statement *"The presence of a tutor was necessary for me to complete this course"*

unlikely that the feedback was directly responsible for this difference. Perhaps the learners have confidence in XeL's ability to replace the tutor by giving timely feedback and accurately measuring their results. This is, of course, highly speculative.

Whatever the reason, these results suggest that XeL may be a good option for remote, self-paced learners. As yet, no remote learners have completed a course with these XeL features enabled (some remote learners used eLGrid before many of the XeL features were developed), but it is hoped that remote learners can be encouraged to use eLGrid in future, and this would give an opportunity to further investigate this interesting effect.

6.5 General attitudes towards eLGrid with XeL

Unchanged attitudes

The general level of learner satisfaction appears to be similar between the XeL and non-XeL groups. A t-test found no significant differences between the responses to any of the following general statements which asked learners about their attitudes to the course materials and the ease-of-use of the eLGrid system

- The course expected too much prior knowledge
- The course was too simple, explaining things that I already knew
- The practical instructions were easy to understand
- The practical environment (WebCom IDE) was easy to access
- The practical exercises helped me to understand and apply the course concepts
- The eLGrid eLearning application is easy to use
- The traffic-light indicators helped me to navigate the course
- In the future I will probably use the technologies taught in this course (WebCom)
- If I need to learn about new Grid technologies in future I would use the eLGrid eLearning system

These results suggest that XeL does not make eLGrid any easier for learner's to use or to make learners *like* the eLearning infrastructure more. Nor does it make any particular difference to their perception of the difficulty of the material presented.

Changed Attitudes

Interestingly, a statistically significant difference was found in the responses to the statement *The theory in this course supported and complemented the practical exercises* between the two groups[6], with the students in the XeL group more likely to disagree with this statement. The effect was stable across the 4BA9 course and both CS4021 courses, as is shown in Figure 6-6.

The reason for this difference is unclear. The theory components were identical for the groups with and without XeL. While the practical was simplified for the XeL group, the simplification focussed mainly on reducing the complexity of the functionality that learners were required to implement within their custom WebCom nodes.

[6] α=0.05, 95% Confidence Interval: [0.109 1.572]

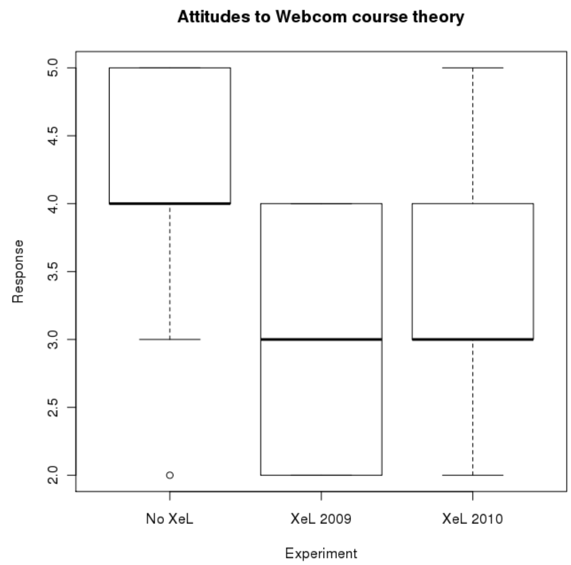

Figure 6-6: Responses to the statement *"The theory in this course supported and complemented the practical exercises"*

This more complicated functionality implemented as part of the 4BA9 practical was not covered by the eLGrid WebCom course, but rather, was covered during the lectures. The aspects of WebCom usage based on material in the eLGrid WebCom course that learners were required to demonstrate remained the same.

One might conjecture that this is simply an artefact of the course delivery method whereby learners had fewer practical sessions with a tutor, and a personal rather than a group project. Alternatively, there may be some real effect of XeL at work here. Perhaps the increased learner focus on the practical exercises makes the theory seem less relevant, or perhaps the fact that learners in the XeL group were more likely to actually work through the practical exercises highlighted deficiencies in the theoretical aspects of the course which other learners more inclined to skip the practicals did not notice. More investigation would be required to better understand the reasons for this difference.

6.6 Attitudes towards integration of the practical exercises

Understanding of the traffic-light indicators

All students were asked to respond to the statement *"I understood how my test answers affected the traffic-light indicators"*. In addition, those students using eLGrid with the XeL integration of the practical exercises were asked to respond to the statement *"I understood how the results of my practical exercises affected the traffic-light indicators"*. The latter question was intended to ascertain whether learners understood how XeL worked and how it extended to the practical exercises. Both questions also give a feel for how intuitive the annotated traffic-light menu system is and how well learners can understand and engage with it. Conversely, a lack of understanding of its function would negate the desired effects in providing a simple and intuitive way for learners to review their progress through the course.

No significant difference was found between the non-XeL group and the XeL group in their response to the first statement.

A significant difference was found, however, in the responses to the statement about *practicals* by the students in the XeL group when compared to the responses to the statement about the *multiple-choice test answers* by the non-XeL group[7]. The statements are similarly worded, and the fact that the XeL group appear to have a better understanding of the practical exercise's effects on the traffic-light navigation (part of XeL) than either group has of the multiple-choice tests' effect may suggest that the effect is more easily discriminated in the case of XeL. This statement does not tell us anything about their satisfaction with the feature, or whether it improves learning, but one of the ideas of the eLGrid adaptive traffic-light navigation menu was that it could enhance active learning by facilitating review and reflection on their own learning. The responses to this statement may suggest that this has, to some extent, been successful.

[7] α=0.05, 95% Confidence Interval: [0.079, 1.611]

Evaluation of knowledge by tests and practical exercises

Responses to the other statements relating to the practical exercises were largely very similar between the groups. The mean responses to the statement *"The tests were a good way for eLGrid to evaluate my knowledge"* were 3.93 and 3.92 for the non-XeL and XeL groups respectively, and the mode of both groups was *"4. Somewhat Agree"*.

The mean responses to the statement *"The practical exercises were a good way for me to evaluate my knowledge"* were 4.14 and 4.08 respectively for the non-XeL and XeL groups, and again, the mode of both groups was *"4. Somewhat Agree"*.

Learners in the XeL group were also asked to respond to the statement *"The practical exercises were a good way for eLGrid to evaluate my knowledge"* in the hopes of learning whether they felt that the practicals were a more accurate measure of their knowledge than the multiple-choice tests. However, responses to this statement were identical to the XeL learners' responses to the previous statement about the tests, suggesting that they did not see the practical exercises as being significantly better. Responses to these statements are shown in Figures 6-7 and 6-8.

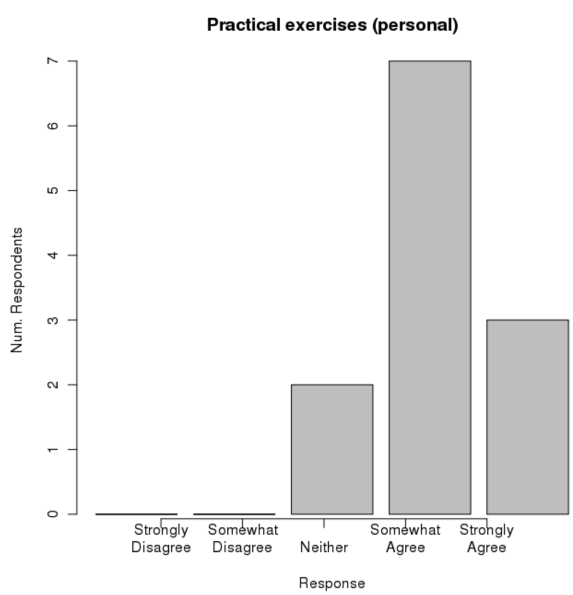

Figure 6-7: Responses to the statement *"The practical exercises were a good way for me to evaluate my knowledge"*

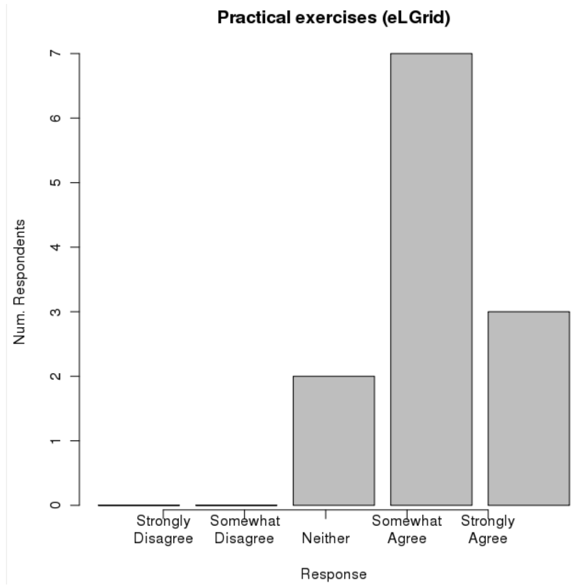

Figure 6-8: Responses to the statement *"The practical exercises were a good way for eLGrid to evaluate my knowledge"*

Usefulness of taking into account tests and practical exercises

To more directly measure attitudes towards the XeL functionality, learners in the XeL group were asked to respond to the statement *"The fact that the traffic-light indicators took into account the results of the practical exercises as well as my test answers was helpful to me"* and this was compared with the non-XeL learners responses to the statement *"It would be useful if the traffic-light indicators took into account the results of my practical exercises as well as my test answers"*, which forecast the existence of XeL.

The majority of learners in the XeL group somewhat agreed that this feature was helpful, while the non-XeL learners were bimodal between *"2. Somewhat Disagree"* and *"5. Strongly Agree"*. The mean of 3.5 and median of 4 for the XeL group is higher than the mean value of 3.286 and median of 3 for the non-XeL group. The responses to the forecasting statement by other non-XeL learners from the first and second evaluation studies were slightly less positive again than the non-XeL 4BA9 group (mean 3.167 and median 3). A t-test failed to find any significant difference in the means of these groups, but the improved rating for this feature is still a positive indication for XeL.

Usefulness of instant feedback

Learners in the XeL group were generally happy with the feedback which eLGrid provided on their practical exercises. Responses to the statement *"The instant feedback on the results of my practical exercises was helpful to me"* had a mode of *"4. Somewhat Agree"* and a mean value of 4.33 (median = 4). These results are shown in Figure 6-9.

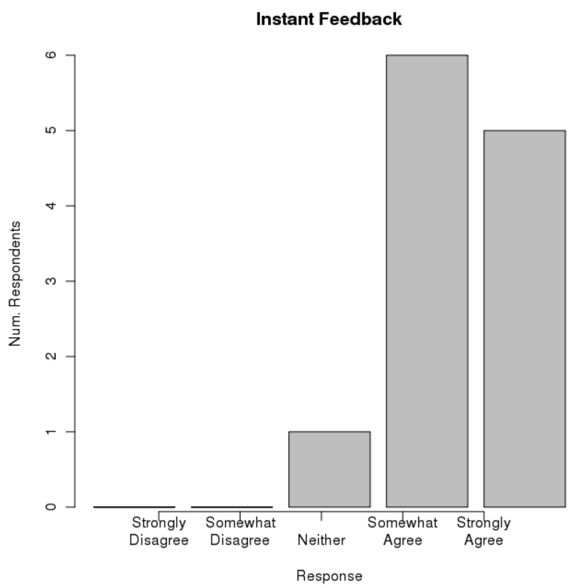

Figure 6-9: Responses to the statement *"The instant feedback on the results of my practical exercises was helpful"*

Close-placement of practical environments and instructions

The integration of the eLearning application with the t-Infrastructure allowed the practical environment for some courses to be presented within the eLGrid eLearning application. This allowed *close placement* of the practical instructions and practical environment.

Learners who took these courses were asked an additional question in order to determine how useful they found this feature. For example, learners taking the *Grid*

Certificates course were asked to respond to the question *Accessing the CA pages from within the eLGrid eLearning system made the practicals easier to complete than if I had had to open the CA in a different window*, while learners taking the P-GRADE and MPI courses were asked to respond to variations of the following statement *"Accessing the P-Grade practical environment directly from within the eLGrid system made the practical easier to complete than where such a direct link was not provided (e.g. when I had to SSH to the Grid UI for the MPI exercises)"*

The responses, with a mode of "4. Somewhat Agree", a mean value of 3.423 and a median of 4, show that the majority of learners agreed that this close placement was helpful, but not very strongly.

Learners taking these courses were also asked the same question as all other learners about ease of access to the practical environment.

A moderate correlation was found between the response to this statement and learners' responses to the statements *"The practical environment was easy to access"* [8] and *"The practical environment is easy to use"* [9].

Thus, one can speculate that the integration of the practical environment within the eLearning application and the consequential close-placement of practical instructions to the environment in which the exercises must be carried out, makes accessing and using the practical environment somewhat easier. No correlation was found between responses to this statement and other statements about the practicals such as *"The practical instructions were easy to understand"*. The correlations are not particularly strong, but they do suggest that this is worthy of further study.

Whether the cognitive load associated with accessing the practical environments when close-placement was employed was actually lower is harder to ascertain. One could look at a comparison of final marks or results of the multiple-choice test questions or XeL practical exercises in order to attempt to measure the difference in learning in groups with and without this close-placement. However, as the only experiments carried out using XeL used the WebCom course which does not employ this close-placement (due to the nature of the WebCom IDE which cannot run within the

[8] Pearson's correlation coefficient 0.35

[9] Pearson's correlation coefficient 0.32

web-browser environment), the data to investigate this effect further is not available.

6.7 Discussion

Accurately measuring the learning achieved from any educational or training activity is not an easy task [203] and in the case of these evaluations, this problem was compounded by the change in course structure between the non-XeL CS4021 course and the XeL 4BA9 course.

There is a clear improvement in course marks in the XeL group, however, and this, taken together with the evidence for increased participation in the practical elements of the course, and the increase in the number of times course concepts related to the practical exercises were viewed, does suggest that the presence of XeL has a beneficial effect on learning.

Although these evaluations took place in the context of a face-to-face course, the difference in responses to the statement "The presence of a tutor was necessary for me to complete the course" between the non-XeL and XeL groups suggests that XeL would be well-suited to remote and self-paced learners, an important cohort of learners that Grid education should cater for. The fact that learners with XeL also have higher participation in the practical aspects of the course lends weight to this idea. If the effect is also present with remote learners, then the increased participation in the practicals should help to increase their learning even in the absence of a tutor. Learners were generally happy with the instant feedback provided by eLGrid on their XeL practical exercises and one might speculate that this would also be of benefit to self-paced learners who are usually in a position of having to work without any feedback on their exercises.

There are reasons to be cautious in the interpretation of the results. In addition to the small sample sizes and non-random selection process for experimental subjects, much of the data is not normally distributed, and while the Welch t-test is relatively robust to non-normality and unequal variances, the nature of the data may affect the results presented here.

For the significance tests an α value of 0.05 was chosen as this is the most widely used α value. For these experiments it might have been quite legitimate to choose 0.10. The aim of the experiments is not to conclusively prove the benefit of XeL, but rather to look for a possible interesting scientific result which merits more investigation. It is not useful to be too stringent and ignore an effect that does exist [204]. The small sample sizes, however, and the non-normal distribution of the data discussed above, mean that the risk of error is larger. The choice of $\alpha = 0.05$ guards somewhat against the chance of a type II error (finding a significant difference where none exists).

Human behaviour is difficult to study, and much of the data presented here relates survey answers, which are inherently subjective [192]. Responses may also be subject to various biases [193, 194]. The data from the eLGrid logs gives a somewhat more objective view of learner behaviour, and the test marks assigned to learners also give a measure of learning. However, although the pre-course questionnaire was able to identify whether learners had previously used any of the technologies covered in the courses, there was no way to discriminate between stronger and weaker students, and no objective baseline measure for learners which could be compared to their final performance in order to measure the learning achieved during the course.

It appears to be difficult to design suitable experiments to clearly resolve the questions of the thesis. The experiments described in this and the previous chapter did not simultaneously include both a control group and an XeL group, rather the same eLGrid course was used (WebCom) with different learners over a number of years, with changes in the practical arrangements between the different groups. This, along with the small sample size and non-random selection of test subjects, means that the results must be interpreted with care.

Ideally, if sufficient numbers of participants can be found, experiments with simultaneous randomised control and XeL groups would be conducted in order to better test the questions of this thesis.

Chapter 7

Conclusions and Future Possibilities

7.1 Conclusions

This thesis has explored the area of Grid eLearning, and looked at how it could be enhanced by the application of adaptivity and personalisation, in particular it has explored how adaptivity can be extended to the practical aspects of Grid eLearning by integrating an adaptive eLearning application with a t-Infrastructure which simulates a production Grid.

The prototype t-Infrastructure and eLearning tools presented in Chapter 4 have demonstrated the feasibility of such an approach. The eLGrid infrastructure involves an eLearning application closely integrated with a fair replica of the Grid-Ireland production infrastructure. This has enabled information about the practical environment and the learner's conduct of practical exercises to be captured by the eLearning tools and passed to the adaptive engine for use in further personalisation of the course.

Several evaluations of this prototype have beem completed, and in spite of small sample sizes and other difficulties, they have highlighted a number of potential benefits. It seems, in fact, that adaptive eLearning and XeL are a very good fit for Grid education, with the potential to enhance learning, increase learner engagement, as well as to open up new possibilities for novel eLearning modalities in the future.

The distributed nature of Grid learners and their asynchronous training demand mean that remote self-paced learning is an ideal training method, and using adaptivity and personalisation suits the heterogeneous nature of Grid learners. The evaluations in Chapter 5 confirmed this heterogeneity in the broad range of learner behaviours, reinforcing the need for an adaptive, personalised approach to eLearning.

The evaluations of XeL in Chapter 6 suggest that it may increase learning, with learners who took an XeL-enabled version of the *WebCom* course obtaining significantly higher marks than those using the non-XeL version. It also appears to increase the suitability of the eLearning tool for learners working without a tutor, which may have implications for remote learners. Learner drop-out in online courses is notoriously high [156], and while it was not directly examined in these evaluations, the indications are promising that XeL might be of benefit here.

Furthermore, XeL has been successfully used in a blended learning scenario with learners using the eLearning environment during face-to-face lab sessions and for self-paced study.

While the majority of learners had generally positive opinions about the practical exercises in the Grid courses delivered via eLGrid, outside of the XeL experimental groups, few learners appeared to spend much time viewing the concepts related to the practical exercises. By contrast, with XeL, learners engaged more with the practical exercises even where they were not part of the final assessment. They may also have been less inclined to use a trial-and-error approach when answering the multiple-choice test questions.

Learners appear to find the provision of an annotated traffic-light menu system useful, and they indicate that it helps them to navigate the courses. Learners also had a good understanding of how their responses to the multiple-choice tests and their practical exercises affected the traffic-light menu system. Whether or not this navigation system facilitates reflection and review, however, is unclear.

The timely feedback provided on the learners' practical exercises was also appreciated by learners, according to their survey responses, but again, it is hard to say to what extent this actually contributed to learning.

There is also some evidence that integration of the eLearning tools and the practical environments makes the practical environment easier to access and use. One might speculate that this could be due to a reduction in extraneous cognitive load associated with switching environment between the eLearning application and the t-Infrastructure, facilitated by the single sign-on and close-placement of practical instructions to the practical environment.

By using a replica of the real production Grid, there is the possibility of actually encouraging the learner to become sidetracked and work on issues that are not relevant to the curriculum. This effect has led some researchers to argue that simulation is, in fact, a better option than giving learners direct access to the real-life scenario [205]. It seems, however, that with careful problem design, and the use of software tools which naturally have some constraints on what learners can and can't do, this is less likely to be a problem. The benefits of the integration in terms of increased learning and increased engagement with the practical elements may well outweigh any such problems, especially in a situation where accurate simulation is not feasible, such as when dealing with a large and complex distributed system such as the Grid.

7.2 Future Possibilities

While the evaluations presented here show the promise of the integrated XeL approach, there are many possible benefits that have not been fully explored. The eLGrid infrastructure, and the XeL approach, open up possibilities for new learning modalities, which might also be further investigated.

First let us consider the eLGrid t-Infrastructure. It is effectively a "virtual Grid", and this should allow packaging in virtualised form for submission as a job to a production Grid or Cloud. More dynamic t-Infrastructure provisioning may be possible using cloud and virtualisation technologies. Tools such as StratusLab or DemoGrid may be used for this purpose. The current virtualised architecture should make such a transition relatively straightforward, but the cloud provisioning model would simplify the process of creating new t-Infrastructure nodes and allow flexible management of training resources to meet peak loads.

It would be useful to open authentication and authorisation so that more undergraduates can access eLGrid, for example, by issuing anonymous short-term certificates. This would serve to get enough numbers to enable a more rigorous double-blind style trial. As a very valuable side-effect, this may help embed t-Infrastructure and Grid into the academic programme, thus allowing, for example, all registered third-level students to get access to the Grid. This is the case in countries such as Greece [139] where students are issued with Grid credentials as part of their registration process, but has the advantage however, that using short-term credentials limits the security exposure where a student might lose or share their credentials.

Work in EMI to standardise authentication and authorisation across four previously separate Grid middleware solutions (ARC, dCache, gLite and UNICORE) will enable the integration of multiple different infrastructures into eLGrid, allowing courses to be run for more than one middleware [97].

eLGrid has so far harnessed only the compute and data features of the Grid, but there is also the possibility of integrating Grid-connected instruments and devices. These can allow learners to practice using these devices rather than using a simulated version or working directly in the lab. The same objective might be achieved with a virtual lab or device, but its effectiveness will depend on how closely it mirrors the real world object or instruments. By using the real devices, which are already connected to the Grid, it is possible to bypass issues of the accuracy of the simulation, and ensure that the learners can practice on a real system. Issues of security, consequences of user-error and scheduling all have to be taken into account, but many universities already have some of these expensive instruments for teaching purposes which could be connected to the Grid and shared among multiple educational institutions with adaptive eLearning courses delivered via eLGrid. Information about learners use of these instruments could then be captured and used for personalisation and assessment.

As for the eLGrid adaptive eLearning application, the software clearly can be improved. For example, as discussed in section 3.4.2, adaptation to the learner's style is common in adaptive eLearning systems. An individual learner's preferred style is generally determined via a pre-course questionnaire. However, the processing power of the Grid opens up the intriguing possibility of using statistical data-mining techniques to analyse a learner's usage of the eLearning system, identify patterns, and automatically add learners to the most appropriate group [206]. Using the power of the Grid would allow such analysis to be done in real-time and learners' profiles to be repeatedly updated as new patterns emerge from the data.

It might also be possible using data mining to identify within the concept space which are core concepts for learners and which are non-core. The non-core concepts could be subject to adaptive navigation and not presented to all learners, while the core concepts would always be presented.

A Grid-enabled version of Weka, the data mining software used in the data analysis for this thesis, has been implemented [207] by researchers from the University of Calabria in Italy and this could be integrated into eLGrid in the same way as eLGrid's t-Infrastructure. The same approach could be taken with integration of the R statistical package. This could allow very sophisticated on-the-fly learner modelling to be achieved.

Of course, a considerable amount of work would be required in order to make this a reality. While the work done in this thesis suggests that it may be a useful approach, the data mining techniques of this thesis would have to be performed repeatedly on different cohorts of learners using the same course in order to see whether stable groupings emerge. Similar analysis would also need to be applied to other courses, ideally with larger number of students to determine whether any techniques which appeared useful for, say the Webcom course, were more broadly applicable. If it were found to be possible to automatically group learners based on statistical or data mining techniques, then further work would be required in order to map suitable personalisation methods to these groupings, and this work is not inconsiderable.

Implementing even the more concrete above improvements is likely to ameliorate any lack of personal support and feedback to a learner (see section 2.1.4). Such improvements might also be reinforced by introducing a "Tutor" at account registration. This would entail embedding some form of personalised communication between them within each page, possibly with notifications and a searchable archive.

To complement XeL's capture of information about the practical environment and exercises, an interesting possibility might be to allow learners to submit the entire code for their programming assignments to eLGrid. This would allow tutors to view and assess the learner's code within the eLGrid system, and this assessment could feed into the personalisation of courses along with the automated marking of the results of the learner's jobs. In the future, automated code-checking tools might be run to produce a simple (though perhaps not particularly accurate) summary of the code quality. This, in combination with the automatic marking of the practical results, could provide a significant assisted-assessment tool for teaching Grid programming, allowing tutors to communicate with learners, view their practical results and code, and potentially to resubmit their Grid jobs for further verification, all from within the eLearning application.

Current trends in active learning emphasise social and collaborative learning activities, and this area has not yet been explored for eLGrid. The addition of a learner forum, wikis, or other collaborative tools would be a simple way to add some collaborative functionality. However, the collaborative power of the Grid may permit this to be further extended. Collaborative tools such as Access Grid could be integrated, allowing learners to communicate using Grid technologies and work on collaborative group projects involving real-time simulations or access to online experimental equipment. This would allow a student group to use the tools locally to rapid-prototype their eLearning of multiple intellectual strands for a group project, and to undertake individual, then group experiments, such as modelling or simulations. The learning content can be further personalised based on the results of these jobs, for example, to present content on parallel algorithms appropriate to the results obtained, to enable further analysis.

In the wider community, and certainly over the long-term and contingent on significant investment and more positive proof of efficacy, the principles presented in this thesis could be extended to school outreach programmes, lab-based projects for undergraduates, MSc/PhD programmes, industrial training, lifelong Continuing Professional Development, etc., even just within a local-area context. Distance group learning, using the tools in a wide-area context, would enable extension to competitive science challenges analogous to those run by business schools, scientific collaborations, exploring potential scientific solutions to specific problems, pre-proposal brainstorming, personalised learning-enriched online-Ads, or even emergency response - in fact anything where wide-area groups must tackle problems that might require rapid (ergo adaptive and executable) eLearning.

As the number of courses and supported communities increases then individual communities may prefer to limit the offered courses to the scope of their community or cognate set of communities. With multiple communities, this obviously maps well onto a multi-tenant Software as a Service (SaaS) scenario, where the client communities run a single instance of the eLGrid adaptive eLearning application, but customised (personalised) with their own courses, and within their own security and database contexts. On a pan-European or global basis, this would enable significantly improved cost/benefit.

Thus the integration of the eLearning environment with Grids and e-Infrastructures benefits the users of these e-Infrastructures. eLearning can be made available to all users of the infrastructure, be they novice learners (for example school or undergraduate students), MSc or PhD students or trained practitioners. The percolation upwards from school to undergraduate to postgraduate to lifelong learning is particularly notable and desirable. The infrastructure which a practitioner uses for their work or research thus becomes the same tool used for their learning and continuing professional development. This possibility is unique to fields which use e-Infrastructures and should be taken advantage of.

Turning to the evaluation experiments, to enable more definite conclusions to be drawn, larger sample sizes are clearly required, along with experiments that simultaneously include both a control group and an XeL group. It may be possible to arrange a "double-blind" trial, for example, by offering an interesting course to all science and engineering students in Trinity College Dublin, where both students and staff would need to be unaware of whether XeL is present.

Despite running controlled experiments over three years, I cannot yet say definitively that XeL does add benefit. It appears to be hard to design the correct experiments to clearly resolve the questions of the thesis. The obvious next step outlined above, a double-blind trial, where both students and staff are unaware of whether they are in the XeL group or the control group, proves more difficult to design without the presence of XeL becoming quite obvious to students and staff alike.

The primary difficulty is that integration of the practical exercises with the eLearning application directly affects aspects of the appearance of the user interface, so it seems unlikely that the value of this integration can be established in a double-blind manner in any learner cohort that is likely to discuss the trial between themselves. It might still be evaluated explicitly with a large and distributed set of decoupled remote learners, where a randomly selected half experienced XeL and the other half not. Let us call this Trial A.

It is likely to be more straightforward to design a double-blind trial to establish the value of the eLearning application taking into account each learner's conduct of the exercises. For example, a double-blind trial could be designed where the practical exercises were integrated with the adaptive eLearning application, but for a randomly selected half of learners the application did not take into account the conduct of the exercises. Apart from the behaviour of the traffic-lights, all learners would experience the same user interface. If the traffic-lights were disabled, or restricted to navigation only then even this effect would be eliminated. Let us call this Trial B.

Only small modifications to eLGrid would be needed to accommodate Trials A and B as these could essentially be considered another form of personalisation with information about which group the learner belongs to (XeL or control) stored in the learners' profile.

Finally, regarding the courses offered by eLGrid, the possibility of automatically creating some of the course components using resources such as Digital Libraries should be further explored. Currently much of the content for the eLGrid courses comes from the ICEAGE Digital Library, however, this content has required substantial modification before it is ready to be used in eLGrid courses. This is because the formats published in Digital Libraries are generally quite coarse-grained materials such as Powerpoint presentations.

The ICEAGE digital library [208], however, has been steadily expanding the types of material that it contains, and it now has a number of items which could potentially be used almost as-is in an eLGrid course. These include practical exercises, video clips from live-training events and fine-grained learning content. The metadata is also improving and it may soon be possible to generate some structural course components automatically from the metadata in the Digital Library. For example, by querying the metadata for all materials in the Digital Library tagged as belonging to a single course or module, it might be possible to generate a partial Concept Space or Narrative model.

A very futuristic possibility is to create workflows that used AI techniques to extract core concepts and related concept spaces from, say, Powerpoint presentations, in a way that automatically repurposed the presentations as eLGrid courses. A spider could then be employed to thereby harvest an extremely large course set from the web, filtered by human moderators. Once again, the processing power of the Grid could enable complicated natural-language processing and concept mapping to be achieved which might not otherwise be feasible. Considerable work would be required before such a possibility could be realised, but a prototype version which produced learning content as part of an authoring tool for tutors and educators might be a useful first step.

7.3 Final Thoughts

This research has had a broad focus, sitting as it does between the fields of eLearning and Grid Computing. I have attempted to bring the best of educational and eLearning theory to bear on the topic of Grid education, and this has entailed a range of topics from fundamental educational theories, through learning theory, to Grid computing and the challenges therein, with a smattering of data gathering and statistical analysis on the way.

I have also had the pleasure to be involved in many Grid education groups and found myself in the position of helping to drive forward related standardisation and recommendations.

Having come to this research with some personal experience of distance learning (I completed my BSc. by distance learning with Dublin City University's *Oscail* distance learning centre), I found the chance to investigate technology assisted learning, with its potential to transform distance education, to be a fascinating endeavour, and I have thoroughly enjoyed the experience.

Appendix A

Implementation Details

This chapter describes the prototype eLGrid which is an implementation of the XeL-enabled Adaptive Grid eLearning architecture described in Chapter 4.

A.1 t-Infrastructure

A t-Infrastructure called *eLGrid* has been implemented for this thesis. The architecture of this solution is shown in Figure A-1.

The t-Infrastructure provides a fair replica of the Grid-Ireland production infrastructure. The replica sites are hosted on virtual machines on a number of physical servers, and Grid services that are shared by all the replica sites are likewise hosted on virtual machines.

The firewall server ensures that the t-Infrastructure is isolated from the production Grid while providing limited connectivity where required. This machine also runs certain necessary services such as the TransDeploy tool and an install server (see later in this section). As the intention is to replicate the production infrastructure as closely as possible, network aliases are configured on the firewall machine corresponding to the network addresses of the site-specific default gateway, install server and other services. This allows the replica Grid site nodes to connect to this machine for these services without changes in their configurations.

The tools used in this solution are described in more detail below.

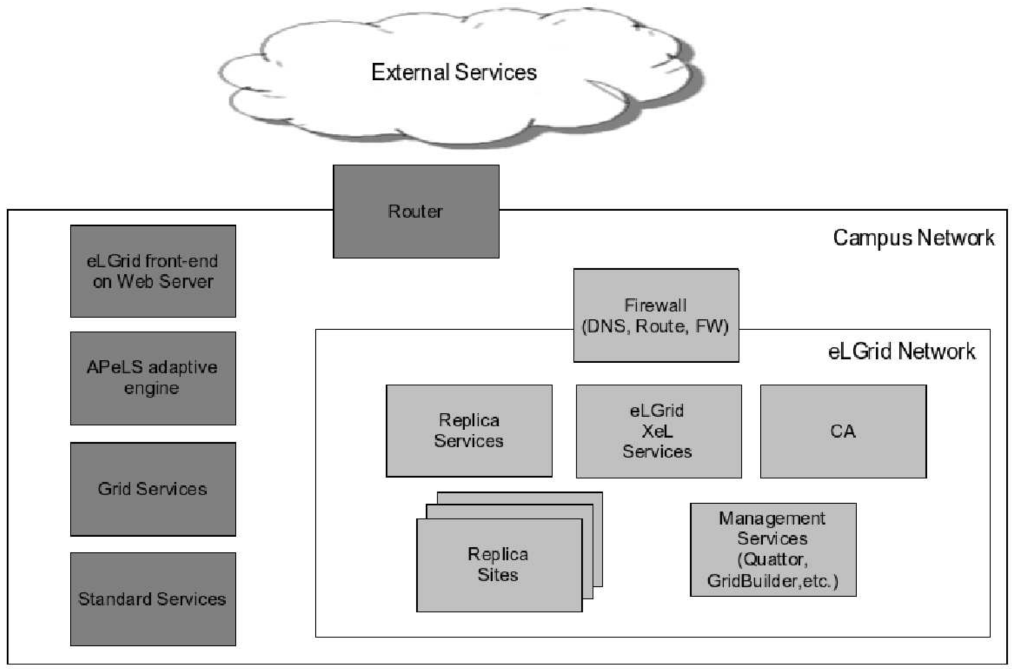

Figure A-1: The eLGrid t-Infrastructure architecture

A.1.1 Virtualisation

The need for any t-Infrastructure to mirror the production infrastructure as closely as possible has already been discussed. The size of the Grid infrastructure and the number of necessary services, however, means that physically replicating even part of the Grid infrastructure would be a huge and costly undertaking.

The use of virtualisation, fortunately, allows replication of a large number of Grid nodes and servers, while significantly reducing hardware requirements. There are a number of virtualisation systems available for X86 architectures which could allow creation of a large virtualised testbed using inexpensive off-the-shelf hardware. The system chosen for this thesis was Xen [209]. The author's host research group has used Xen since 2004 for its testing and development sites [210] and also for the Grid-Ireland production infrastructure [211].

Xen was chosen not simply because it is already used in the Grid-Ireland infrastructure, but also because at the time of adoption it had a price/performance edge over the other virtualisation systems which are available for the x86 platform. Tests conducted by Childs *et al.* [211], and confirmed by Hardt and Berlich [212] found that Xen gave 90-95% of the performance of a reference Grid machine without Virtualisation, outperforming alternatives such as User Mode Linux, VMWare, etc. Exact performance varies depending on the hardware available and the type of job being run, but the performance of Xen is certainly quite adequate for training needs.

Thanks to the foresight of the HEA PRTLI3 CosmoGrid project (led by Prof. Luke Drury, DIAS), eight Dell dual-processor dual-core machines, each with 8GB of memory and 500GB of disk were procured for this purpose, and installed with a Xen hypervisor and Linux kernel. Each replica site consists of a minimum of five virtual machines: a User Interface (UI), a Compute Element (CE), a Storage Element (SE), a Grid monitoring test node, and at least one Worker Node (WN). The number of worker nodes can easily be increased to handle greater load, for example when a course is scheduled to take place.

Allowing approximately 400 MB of memory and 8 GB of disk space for the privileged domain 0 host operating system (Dom0) and guest VMs (some additional disk space is used for the Xen base images which are stored in their own LVM partitions) allows replication of up to four Grid-Ireland sites on each physical machine.

This yielded the ability to replicate the entire Grid software and services of the Irish Grid infrastructure, albeit with limited numbers of worker nodes and reduced performance. Nonetheless, the replica is quite scalable. Additional memory would allow more worker nodes. More machines would allow more sites. More modern multicore machines would improve performance.

A.1.2 Certificate Authority

The eLGrid t-Infrastructure has its own Certificate Authority (CA) which issues certificates valid only on the eLGrid infrastructure. The eLGrid CA uses OpenCA [213], the same software as is used by the Grid-Ireland production CA, but it is configured to issue user certificates with a shorter lifetime (by default one month) than those issued by the production CA.

OpenCA can also issue host certificates, revoke certificates, manage certificate revocation lists (CRLs), etc. The CA front end is shown in Figure A-2.

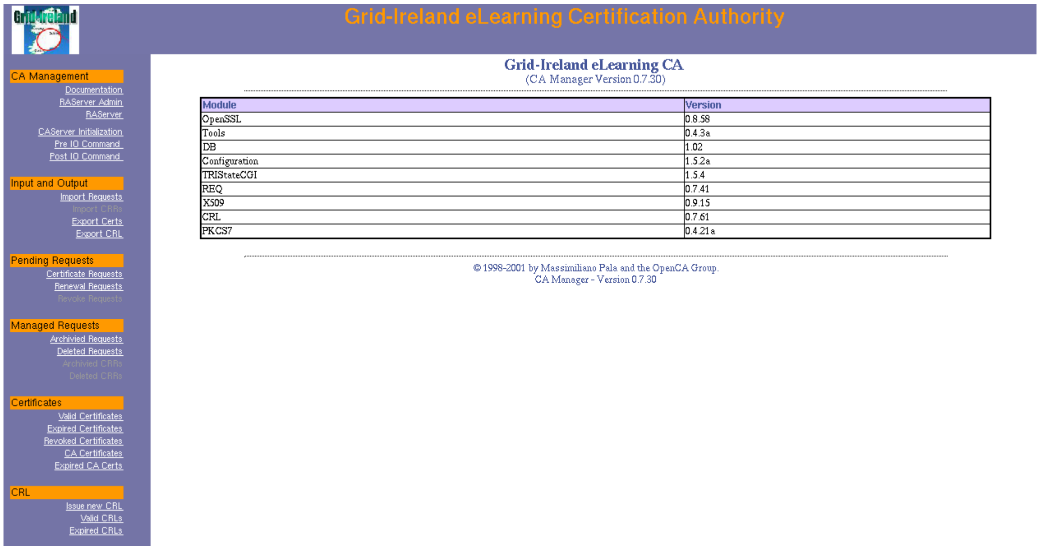

Figure A-2: The eLGrid Certificate Authority front end.

A.2 The eLGrid Adaptive eLearning Application

The eLGrid eLearning application offers personalised courses to learners and integrates with the eLGrid t-Infrastructure. It uses the Adaptive Personalised eLearning Service (APeLS) [214] developed by Dr. Owen Conlan of the Knowledge and Data Engineering Group (KDEG) in Trinity College Dublin [215]. The APeLS service handles manipulation of the learner and course models on which the personalisation of the course is based, while the eLGrid eLearning application itself handles interaction with the learner, displaying the course components, logging, reporting and administration functions.

The eLGrid eLearning application was developed to provide adaptive eLearning courses which are integrated with the eLGrid t-Infrastructure. The application is written in Java and uses the Apache Tomcat Servlet Container [216].

The APeLS engine provides an API for manipulation of the learner and other models on which the course model is based. On its own, however, it is not a complete course delivery system. Rather APeLS provides a set of classes which are used by a course delivery system to handle the personalisation aspects of the course delivery, and the eLearning application handles interaction with the APeLS classes as well as the remaining functionality.

The eLGrid eLearning application and the APeLS engine are both installed on the Grid-Ireland webserver. The original intention was to install these on the firewall machine itself, however, the front-end system must be installed on a system which is accessible to the outside world so that a wide range of learners can access the courses. The Grid-Ireland webserver is accessible on the Internet and was already running a secure Apache webserver [217], thus it seemed a good location on which to install the application. As this webserver is not directly connected to the eLGrid t-Infrastructure, connectivity to the t-Infrastructure is provided by an SSH tunnel through the eLGrid firewall machine.

The Grid-Ireland webserver is secured with X.509 Grid certificates using GridSite [218]. This means that it is possible to only allow users possessing a valid certificate to gain access to the courses. In addition, it is possible for the GridSite software to pass the user's distinguished name (DN) to APeLS as the user name, allowing single-sign-on. The eLGrid user interface is accessible using either a username and password or via single-sign-on with Grid certificates. This ensures that learners who have not yet obtained a Grid certificate can access the system and use the *Grid Certificates* course. The configuration of the single-sign-on for the user interface is described in section 4.5.2.

The eLGrid graphical user interface is a Java Server Pages application running in an Apache Tomcat server. A number of JSP files handle user requests and the main functionality is provided by a set of Java class files.

The design has followed published design guidelines where such exist, for example the *Research-Based Web Design and Usability Guidelines* [219]. In addition to the general layout and design of the application, the instructional content has also been designed to minimise extraneous cognitive load [18] so that it does not militate against learning.

Users generally prefer a consistent User Interface throughout an application. As it is anticipated that most learners will find the eLGrid system via the Grid-Ireland website, the Grid-Ireland look-and-feel has been maintained by reusing the Grid-Ireland header and footer files and CSS files.

However, it is important that eLGrid should also be stand-alone. It should not require the user to be a member of Grid-Ireland, and should not confuse learners by including Grid-Ireland navigation options among the course menus. The Grid-Ireland page-layout is therefore reused, but the Grid-Ireland menus are not displayed, although a link back to the Grid-Ireland website is provided. Furthermore these files are included in such a way as to be easily replaced by alternative header, footer and CSS files if, for example, it becomes necessary to install the system elsewhere. The exception to this is the login screen which is considered part of the Grid-Ireland site and which includes the Grid-Ireland navigation menu.

Once logged-in the learner is presented with a course list and once a course is chosen there are two possible ways for learners to navigate through the course concepts; *Next* and *Previous* buttons and a menu annotated using a *Traffic-light* metaphor.

A.2.1 Next and Previous buttons - Direct Guidance

The *Next* and *Previous* buttons in the eLGrid eLearning system use the principle of Direct Guidance [14] to help the learner navigate through the course content.

The *Next* button recommends the best resource for the learner to view next. This may not always be the next item in the menu. If the learner has already viewed and passed a test in the next concept, for example, then the *Next* button will skip over that concept and bring the learner straight to the next untested concept. Similarly the *Previous* button skips over tested concepts. These tested concepts are still available via the menu, however, should the learner wish to view them again.

A.2.2 Menu - Link Annotation and the Traffic-light metaphor

A course menu is displayed on the left-hand side of the page. The menu is created adaptively and uses link annotation with a traffic-light metaphor [14].

Menu items have a coloured icon placed beside them which provides the user with additional information about the concept linked to each item:

(a) A green icon indicates that the learner is ready to view a concept. It is also used beside course sections to indicate that the learner can view that section.

(b) An amber icon indicates that the learner has already viewed a concept and that they are now ready to be tested in this concept if they so choose. If no test exists in the course then the amber icon will not be used, instead the concept will be marked with a lilac icon (see (d) below) once it has been viewed. If a test is later added then the traffic light colour will change the next time the learner visits the course.

(c) A red icon indicates that the learner has not yet viewed or been tested for the prerequisite materials for this concept and is thus probably not ready to view this material.

(d) A lilac icon indicates that the learner has been successfully tested for this concept, or that they have viewed the material if no test exists for the concept.

(e) A grey icon indicates that the associated course section is not a necessary part of the course for that learner. It might be used in a case where a course is used for students from different disciplines, for example, if a particular section is only required by some of these students.

The eLGrid front-end is shown in Figure A-3, where the *traffic-light* annotated menu is visible on the left of the screen.

A.3 The Adaptive Engine - APeLS

eLGrid uses the Adaptive Personalised eLearning Service (APeLS) is a web-services-based Adaptive Hypermedia System (AHS). This creates personalised courses at run time by adapting the content or the navigation based on the interaction of a number

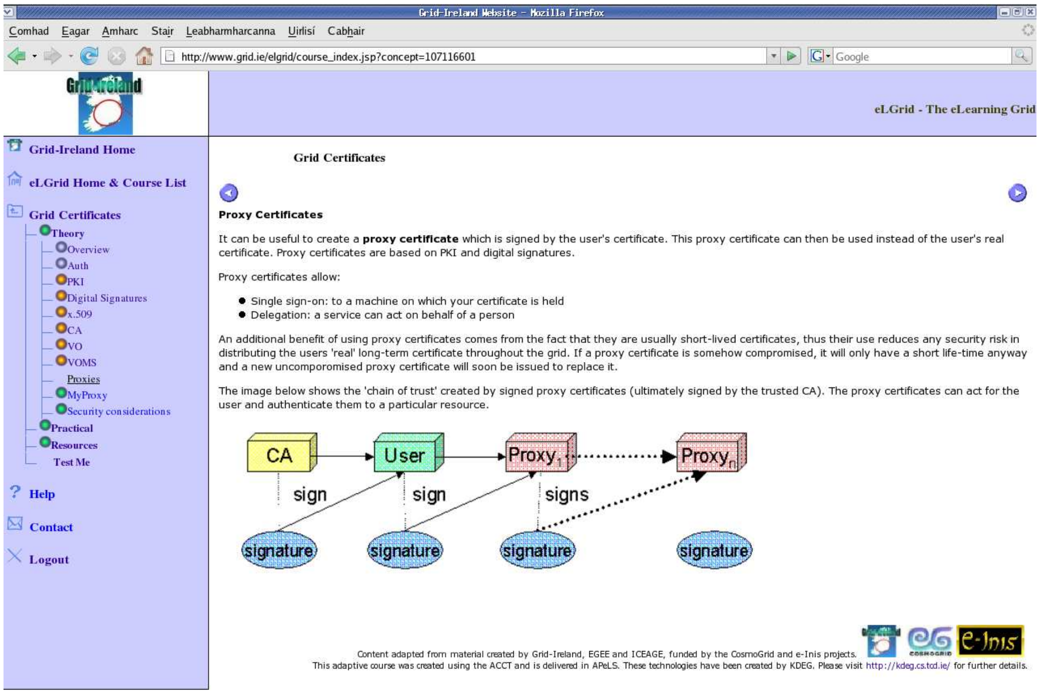

Figure A-3: The eLGrid front-end

of *models* such as *Pedagogic Activity Sequence* (narrative), *Subject Area* (concept space), *Candidate Learning Resources* (learning content), Learner Model, etc.

The models are stored as XML metadata in an eXist XML database [220] and combined at run-time by an AdaptiveEngine Java class to create a personalised course model. A set of XSLT transforms then convert the XML course model into html for display within an eLearning application. The eLearning application may rebuild the course model at any time by calling the appropriate AdaptiveEngine method.

Unlike many other adaptive eLearning systems, APeLS is a generic *multi-model* personalisation engine, i.e. it can personalise a course based on any number of attributes, or in APeLS terminology *"adaptive axes"*. This makes it a flexible tool which can be used to add new personalisation options to a course without rewriting an entire eLearning application.

The system has been used in Trinity College Dublin for delivery of an undergraduate SQL course since 2000 [64].

A.4 XeL Modules

Three XeL modules have been developed at this time. Two of these are simple add-ons or wrappers for a learner's job, which the learner must include in their job when running. The third is an attempt to make a more comprehensive solution involving a daemon which runs within the t-Infrastructure and queries the standard Grid services.

A.4.1 Webcom XeL Node

The Webcom XeL Node is a Webcom graph node which must be included as the final node in a learner's Webcom graph. It takes the output of the learner's graph and sends it via http to the eLGrid application. The application then calls a handler script associated with the practical exercise in question, which decides how to deal with the output. The currently implemented handler script simply stores the data and checks whether it is correct. A response is then sent to the Webcome XeL node indicating whether or not the learner has got the correct answer.

This is the method which was used in the experiments using XeL.

A.4.2 gLite wrapper script

The gLite wrapper script is a simple shell script which takes the name of the learner's executable and the parameters to be passed to it, and then calls that executable script with the parameters. On completion of the learner's job, the output is captured by the wrapper script and sent via http to the eLGrid application. The application then calls a handler script as for the Webcom node.

This solution requires that the learner submit the wrapper script along with their job, although the possibility of modifying the gLite Grid software to automatically call this wrapper script was investigated.

A.4.3 gLite daemon

The third solution uses a daemon which must be installed within the t-Infrastructure. The daemon is sent the learner's DN from their Grid certificate by the eLGrid application. It can then query the status of jobs for that learner via the Logging and Bookkeeping (L&B) daemon. This solution seems the most robust method and has the least impact for the learner. However, frequent changes in Grid middleware coupled with limitations in the types of information available via the L&B service make this difficult to implement in practice. A new version of this daemon is now in development which queries the gLite Workload Management System (WMS) as well as the L&B service to allow it to access more information about the learner's jobs.

Appendix B

The Networking Layer

Each replica VM in the eLGrid t-Infrastructure uses the same public IP address as that used by that machine's counterpart in the production Grid-Ireland infrastructure. This allows identical configurations to be used for the eLGrid t-Infrastructure and the Grid-Ireland infrastructure.

Because of this replication, it was necessary to completely firewall off the t-Infrastructure from the rest of the network. The eLGrid firewall is a dual-homed machine with an 'internal' network device connected to the eLGrid t-Infrastructure and an 'external' device connected to the real production network. Iptables and arptables firewalling and NAT filtering ensure that no traffic from the internal interface can accidentally be sent out onto the external network and thereby to the production network (and vice-versa).

A caching DNS server installed on the firewall machines handles DNS for the physical eLGrid machines (the Xen hosts) and queries the Grid-Ireland production DNS server for the IP addresses of Grid nodes within the production infrastructure. These addresses are then cached for subsequent use, allowing the replica virtual machines to be instantiated with the same IP addresses as their counterparts in the production infrastructure.

On the firewall machine an alias is created for the default network gateway (as distinct from the "Grid gateway") for each site. All packets coming from replicated site nodes within the t-Infrastructure will be directed to the appropriate alias on the internal network interface of the firewall machine. Routing is configured on the firewall to ensure that internal packets coming in on these aliases are redirected back out on the same alias.

IPtables is used for the firewall itself. This has a number of rules designed to ensure that no packets destined for internal replica machines can get out onto the production network. Certain types of packet are allowed through (for example, NTP requests are routed out to the 'real' NTP server) to avoid having to maintain separate instances of all non-Grid services. The firewall rules ensure that even if the replica site's routing is set up incorrectly, they cannot accidentally send t-Infrastructure data out into the real Grid.

Some other services are hosted on the eLGrid firewall machine rather than being replicated as VMs. The replica hosts, however, know only of the real IP addresses for these services. NAT is used on the eLGrid firewall to redirect packets directed at servers which do not exist in the eLGrid network to the firewall machine where these services are hosted.

This reuse of the real Grid-Ireland IP addresses yields an identical configuration to that of the production infrastructure. This is a tried and tested configuration as the same setup is used for the Grid-Ireland testing and development infrastructure [210].

Using an identical setup means that the same configuration files can be used as those used to configure the production infrastructure, thus management and administration of the eLGrid t-Infrastructure can be done using the same fabric management system which is used in production, with changes to configuration files only having to be made and tested once.

Appendix C

Management Services

C.1 The GridBuilder Tool

The GridBuilder [171] tool was developed by Dr. Stephen Childs in the Computer Architecture and Grid Research Group in Trinity College. It provides an easy-to-use web-based user interface which allows the user to quickly configure and start new VMs for Grid infrastructures. The GridBuilder interface is shown in Figure C-1.

Gridbuilder stores a library of filesystem images for standard node types, each as an independent LVM partition. Examples of these standard filesystem images include gridmon (the grid monitoring test node), gridui (the site-specific user entry point (UI)), gridstore (site storage (SE)) and gridgate (site Grid service entry point (CE)).

When a new virtual machine is created, GridBuilder creates a copy-on-write clone of the appropriate LVM partition. The copy-on-write feature of LVM allows clone volumes to be created which only save changes to the base volume, thus saving disk-space. It also speeds up the cloning process.

GridBuilder then mounts the new volume and modifies the configuration. The required configuration is downloaded from the fabric management system (Quattor, Yaim or LCFG).

Once the filesystem volume is unmounted the Xen image is ready to boot. Grid-Builder then boots a Xen VM from this image and any final updates are performed

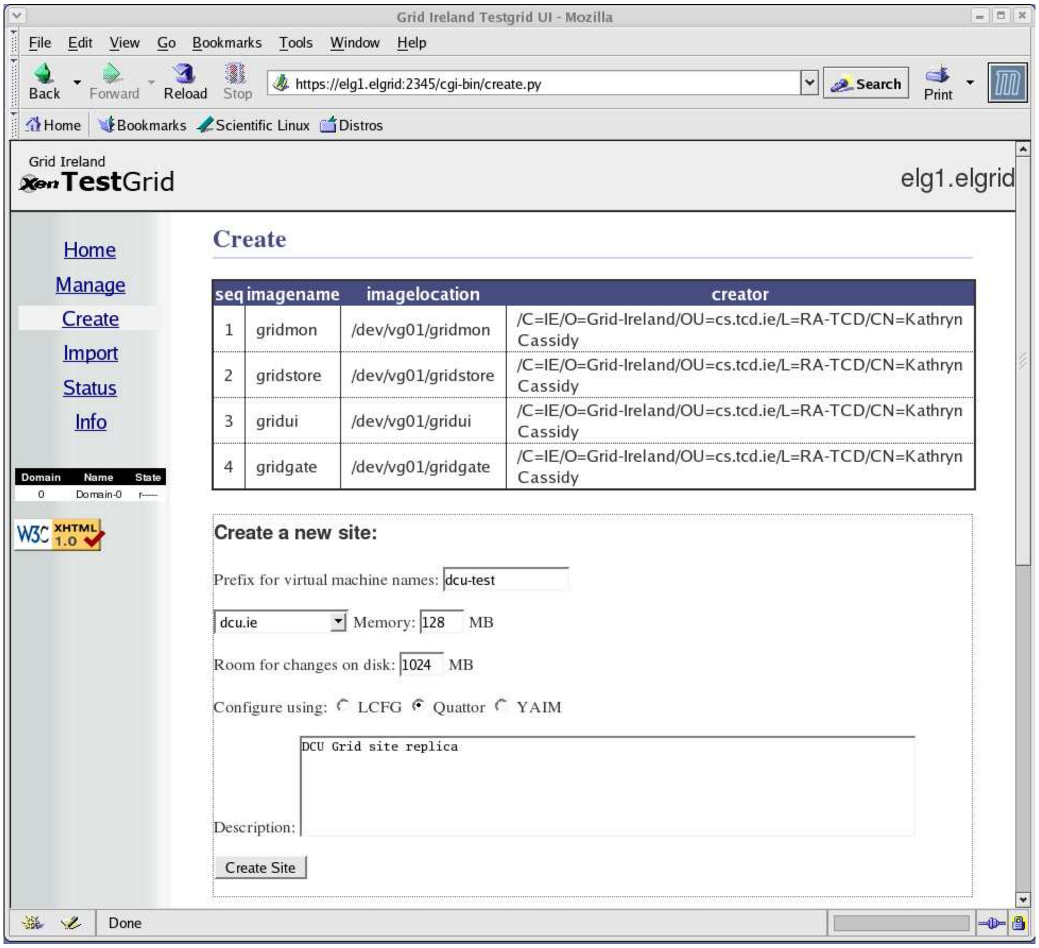

Figure C-1: The GridBuilder Tool allows fast replica VM creation.

by the fabric management system on the VM.

Grid Builder is used in eLGrid to provision the base VMs for the t-Infrastructure.

C.2 Configuration and Management

Any network with more than a few nodes should have a fabric management system to configure, update and install machines, or the workload of system administration will quickly become unscalable. In Grid systems this is particularly important as the sites involved can be particularly large.

In Grid-Ireland the Quattor fabric management system [221] is used. The configurations are stored on each site's install server. Quattor's update mechanisms take over from there, i.e. autonomous nodes running the Quattor client components pull their configurations and update themselves. This autonomy makes the system scalable.

The configuration is stored in templates and sets of preconfigured templates for Grid nodes are available from CERN. The hierarchical structure of Quattor templates means that templates can be combined, for example, a general national Grid-Ireland template can be overridden with a specific site template, which itself can be overridden with a machine-specific template. These templates are stored in a centralised CVS repository on the Grid-Ireland repository server at the Grid-Ireland OpsCentre and deployed to each site's install server.

As previously mentioned the same configuration files are used for the production Grid and for the t-Infrastructure. In the production infrastructure each site has a local install server. The eLGrid t-Infrastructure has just one instance to serve all its replica sites, such that each replica site shares a *"virtual install server"*. In fact this instance runs on the eLGrid firewall, and the virtual install server consists of the web-based configuration repository and the network aliases for the site-specific install servers which can be accessed via the replica nodes for each site. There is also an alias for an eLGrid install server which serves as the install server for the physical eLGrid machines. This install server accesses the Quattor templates on the central Grid-Ireland repository server via an NFS mount. The CVS root is mounted and a read-only CVS server is run locally on the eLGrid install server. This allows profiles to be checked out from this server, but changes cannot be checked in.

Both the physical eLGrid machines and the replica VMs are configured via Quattor. The physical machines require that a new Quattor site be configured and templates created for each machine. However, the configuration of the replica VMs does not require any work as it is possible to use the same configurations as those used to manage the production infrastructure. The Quattor profiles must be available via http on the eLGrid install server aliases in order that they may be accessed by the eLGrid nodes. This is the first step in the deployment process.

C.3 Deployment

In the production Grid, configuration changes to the Quattor configuration templates are copied to the relevant Grid-Ireland production site install servers. To do this Grid-Ireland uses a deployment tool called TransDeploy [172], developed by Dr. Brian Coghlan and Dr. Geoff Quigley in the Computer Architecture and Grid Research Group in Trinity College, and then significantly enhanced by Dr. Stephen Childs (same group).

The TransDeploy tool aims to ensure consistency of sites and to minimise downtime due to upgrades. It also facilitates a push-button deployment process. Upgrades are split into a variable-duration prepare phase and a short-duration upgrade phase (i.e. a two-phase commit, where the commit operation is a simple and fast switch of directory links). Thus overall the upgrade is performed transactionally, in that the entire upgrade is treated as one atomic operation which either fully succeeds or never occurs. And once done, the upgrade can be rolled back simply by reversing the commit.

The prepare phase compiles and tests the profiles and identifies most configuration errors. If the configuration is valid then TransDeploy will proceed to push the configuration templates out to the appropriate site install servers, by instructing those servers to check out the new templates. If this succeeds then the upgrade phase will be executed. Once these valid new configurations are available on the site install servers, the site nodes running Quattor's client components can pull their configurations and update themselves. This may take some time.

There is one TransDeploy instance for the entire Grid-Ireland production network which pushes templates out to the install servers in each site. A separate TransDeploy instance is configured for the eLGrid t-Infrastructure. This is installed on the t-Infrastructure firewall machine and pushes templates to the eLGrid install server aliases within the t-Infrastructure. In principle this could be merged with the production TransDeploy instance to simplify management, but for practical reasons this would require additional development.

The TransDeploy interface for the eLGrid t-Infrastructure is shown in Figure C-2.

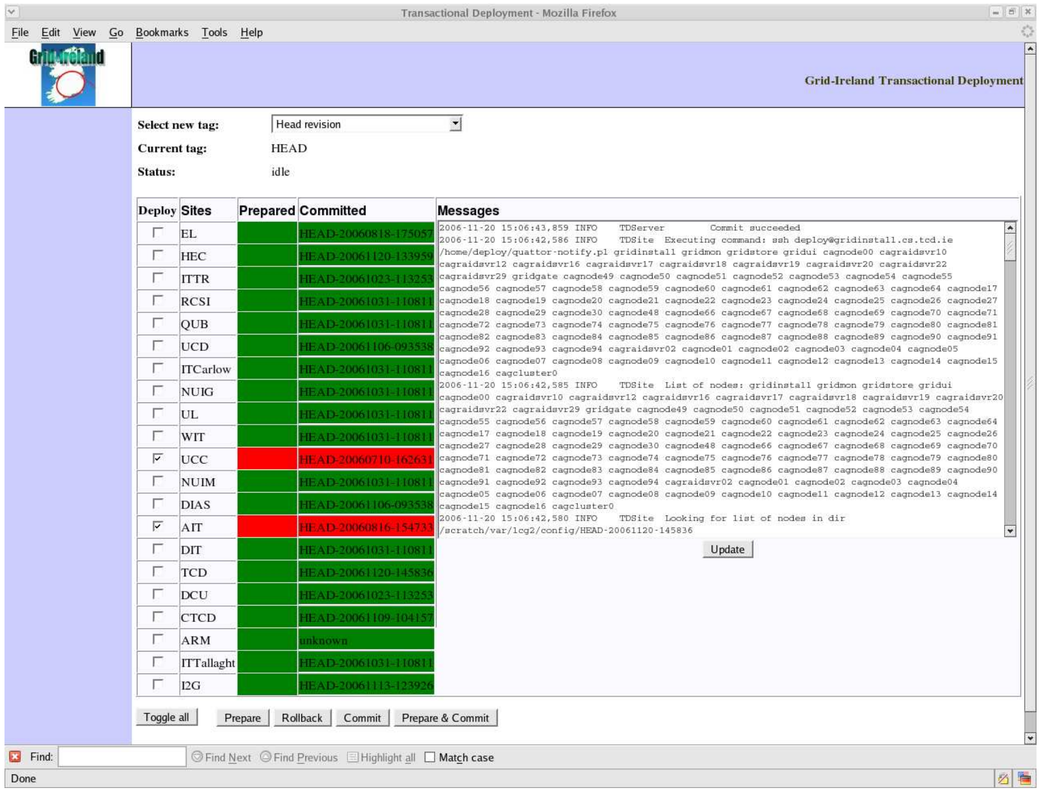

Figure C-2: The TransDeploy tool showing the eLGrid servers.

C.4 Course Construction

The primary tool used to create eLGrid courses is the Adaptive Course Construction Toolkit (ACCT) [173] developed by Dr. Declan Dagger of the Knowledge and Data Engineering Research Group in Trinity College. This provides a graphical user interface for the creation of the major components of an adaptive course:

- *Subject Matter Concept Space*: an ontology of the subject matter including the concepts and their relationships to one another
- *Narrative*: a description of the semantics of the pedagogical strategy of the course. The narrative contains information about the concepts and their *candidate learning resources* (the content which could fulfil the learning goals for

each concept), along with *adaptive axes* or descriptions of the personalisation options for the course.

Courses created in ACCT can then be delivered via APeLS.

ACCT is designed to facilitate the creation of pedagogically sound courses. While not enforced, the tool encourages the course designer to focus on the pedagogical strategy, and to understand the pedagogical basis for the design decisions which they take [173]. This is in contrast to other tools such as common LMSs (WebCT, Moodle, etc.) which focus more the technical aspects with little support for non-traditional pedagogic approaches [222]. Figure C-3 shows the ACCT user interface for creating a concept space model.

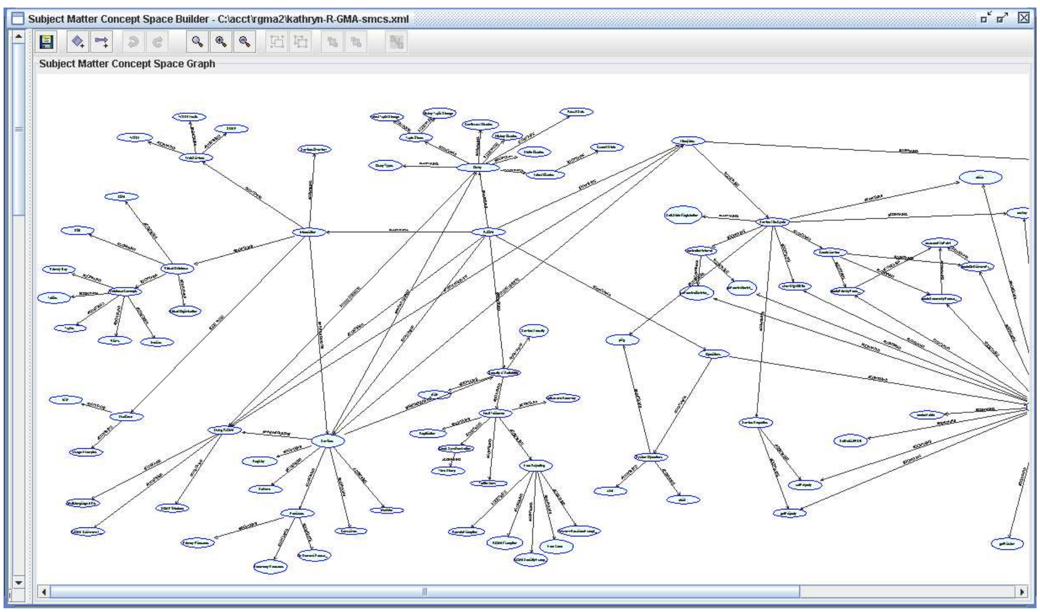

Figure C-3: Creating a concept space in ACCT

Other authoring tools were used where necessary. For example, understanding the structure of a very large concept space from a traditional concept map is not always straightforward, in particular where the job of creating the course is shared among multiple tutors. In this case, additional tools are beneficial.

A transform script was created to convert from the ACCT Concept Space format into a hyperbolic tree format which could be viewed using the HyperGraph applet [223]. The hyperbolic tree representation can allow easier comprehension of large concept spaces because the whole concept space (the 'big-picture') can be represented on-screen in a compressed form while a particular region is brought into focus and thus enlarged. This allows the relationships between even distant nodes to be viewed while at the same time making out the detail of particular nodes of interest [224]. Figure C-4 shows the hypergraph for the concept space of the R-GMA course (see 4.12.3).

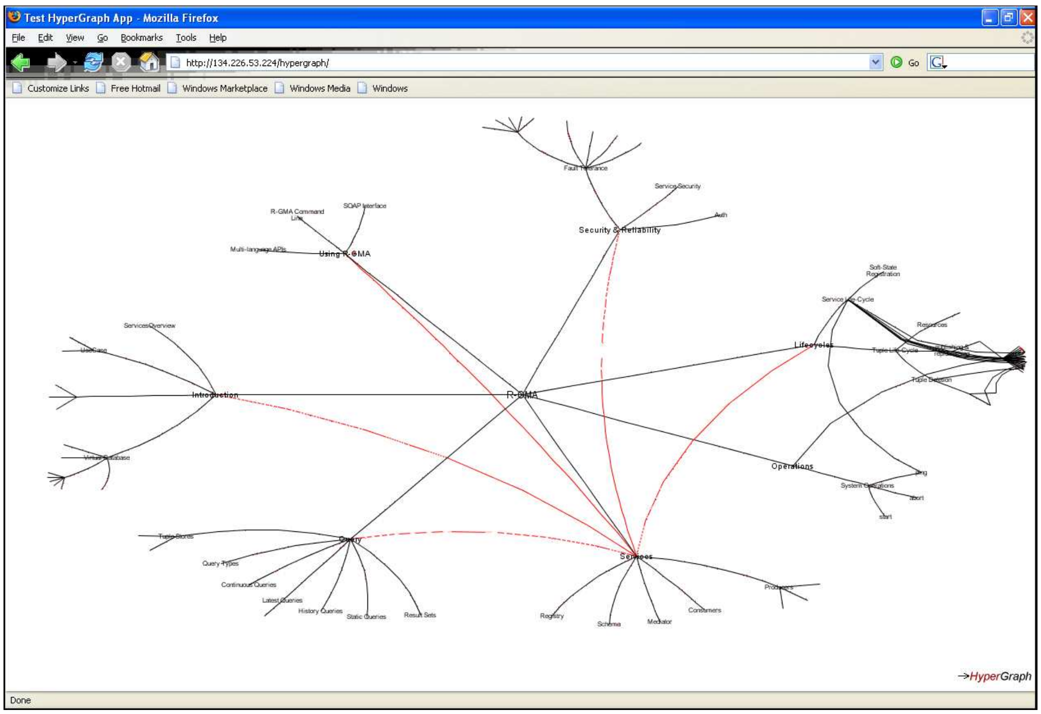

Figure C-4: Hypergraph representation of the R-GMA concept space

Appendix D

eLGrid Learner Model

eLGrid uses a learner model metadata scheme based on the default schema used by APeLS, but extended to include some Grid-specific elements such as VO and discipline.

Two main types of adaptivity are included, concepts.viewed and concepts.learned. Concepts.viewed is used to list the concepts that the learner has viewed, while concepts.learned lists only the concepts in which the learner has demonstrated competency. Competency can be demonstrated by completing a test for that concept via the *Test Me!* feature of eLGrid or by completing a related XeL practical exercise for that concept which automatically uploads the practical results to eLGrid.

The learner model shown here is used in the creation of the eLGrid annotated traffic-light menu system.

```
<?xml version="1.0" encoding="UTF-8"?>
<learner>
    <general>
        <identifier/>
        <name>
            <surname>Cassidy</surname>
            <forename>Kathryn</surname>
        </name>
        <discipline>Computer Science</discipline>
<vo>cosmo</vo>
```

```
    </general>
    <educational>
        <adaptivity>
            <adaptivitytype name="concepts.learned">
                <set type="ALL">
                    <candidate>
                        <langstring lang="en">330185283</langstring>
                    </candidate>
                </set>
            </adaptivitytype>
        </adaptivity>
    </educational>
    <navigational>
        <adaptivity>
            <adaptivitytype name="concepts.viewed">
                <set type="ALL">
                    <candidate>
                        <langstring lang="en">329954512</langstring>
                    </candidate>
                </set>
            </adaptivitytype>
        </adaptivity>
    </navigational>
    <navigational>
        <adaptivity>
            <adaptivitytype name="concepts.viewed">
                <set type="ALL">
                    <candidate>
                        <langstring lang="en">330185283</langstring>
                    </candidate>
                </set>
            </adaptivitytype>
        </adaptivity>
    </navigational>
</learner>
```

Appendix E

eLGrid Content Model

This section shows an example metadata file conforming to the eLGrid schema. The example chosen is a file for the *Database Indices* topic of the eLGrid *R-GMA* course.

```
<?xml version="1.0" encoding="UTF-8"?>

<lom xmlns="http://www.imsglobal.org/xsd/imsmd_v1p2"
     xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
     xsi:schemaLocation="http://www.imsglobal.org/xsd/imsmd_v1p2 imsmd_v1p2p2.xsd">

<general>
  <title>
    <langstring xml:lang="en">Database Indices</langstring>
  </title>
  <catalogentry>
    <catalog>URI</catalog>
    <entry>
      <langstring xml:lang="en">
        http://blacktower.cs.tcd.ie/rgma_course/indices.html
      </langstring>
    </entry>
  </catalogentry>
  <language>en</language>
  <description>
    <langstring xml:lang="en">Database indices in R-GMA</langstring>
```

```
  &lt;/description&gt;
  &lt;keyword&gt;
    &lt;langstring xml:lang="en"&gt;R-GMA&lt;/langstring&gt;
  &lt;/keyword&gt;
  &lt;keyword&gt;
    &lt;langstring xml:lang="en"&gt;Database&lt;/langstring&gt;
  &lt;/keyword&gt;
  &lt;keyword&gt;
    &lt;langstring xml:lang="en"&gt;Indices&lt;/langstring&gt;
  &lt;/keyword&gt;
  &lt;keyword&gt;
    &lt;langstring xml:lang="en"&gt;Index&lt;/langstring&gt;
  &lt;/keyword&gt;
  &lt;keyword&gt;
    &lt;langstring xml:lang="en"&gt;Indexes&lt;/langstring&gt;
  &lt;/keyword&gt;
  &lt;coverage&gt;
    &lt;langstring xml:lang="en"&gt;Republic of Ireland&lt;/langstring&gt;
  &lt;/coverage&gt;
  &lt;coverage&gt;
    &lt;langstring xml:lang="ga"&gt;Éire&lt;/langstring&gt;
  &lt;/coverage&gt;
  &lt;coverage&gt;
    &lt;langstring xml:lang="en"&gt;Europe&lt;/langstring&gt;
  &lt;/coverage&gt;
  &lt;structure&gt;
    &lt;source&gt;LOMv1.0&lt;/source&gt;
    &lt;value&gt;atomic&lt;/value&gt;
  &lt;/structure&gt;
  &lt;aggregationlevel&gt;
    &lt;source&gt;LOMv1.0&lt;/source&gt;
    &lt;value&gt;1&lt;/value&gt;
  &lt;/aggregationlevel&gt;
&lt;/general&gt;

&lt;lifecycle&gt;
  &lt;version&gt;
```

```
      <langstring>1</langstring>
    </version>
    <status>
      <source>LOMv1.0</source>
      <value>draft</value>
    </status>
    <contribute>
      <role>
        <source>
          <langstring xml:lang="x-none">LOMv1.0</langstring>
        </source>
        <value>
          <langstring xml:lang="x-none">Publisher</langstring>
        </value>
      </role>
      <centity>
        <vcard>
BEGIN: vCard
ORG: COSMO
END: vCard
</vcard>
      </centity>
      <date>
        <datetime>2007-04-17</datetime>
      </date>
    </contribute>
</lifecycle>

<metametadata>
    <catalogentry>
      <catalog>URI</catalog>
      <entry>
        <langstring xml:lang="x-none">
          http://blacktower.cs.tcd.ie/metadata/rgma/indices.xml
        </langstring>
      </entry>
    </catalogentry>
```

```
  <contribute>
    <role>
      <source>
        <langstring xml:lang="x-none">LOMv1.0</langstring>
      </source>
      <value>
        <langstring xml:lang="x-none">Creator</langstring>
      </value>
    </role>
    <centity>
      <vcard>
BEGIN:vCard
Cassidy;Kathryn
END:vCard
</vcard>
    </centity>
    <date>
      <datetime>2007-04-17</datetime>
    </date>
  </contribute>
  <metadatascheme>CanCore 1.0</metadatascheme>
  <metadatascheme>IEEELOM:1.0</metadatascheme>
  <language>en</language>
</metametadata>

<technical>
  <format>text/html</format>
  <size>635</size>
  <location type="URI">
    http://blacktower.cs.tcd.ie/rgma_course/indices.html
  </location>
  <requirement>
    <type>
      <source>
        <langstring xml:lang="x-none">LOMv1.0</langstring>
      </source>
      <value>
```

```
        <langstring xml:lang="x-none">Operating System</langstring>
      </value>
    </type>
    <name>
      <source>
        <langstring xml:lang="x-none">LOMv1.0</langstring>
      </source>
      <value>
        <langstring xml:lang="x-none">none</langstring>
      </value>
    </name>
  </requirement>
  <requirement>
    <type>
      <source>
        <langstring xml:lang="x-none">LOMv1.0</langstring>
      </source>
      <value>
        <langstring xml:lang="x-none">Browser</langstring>
      </value>
    </type>
    <name>
      <source>
        <langstring xml:lang="x-none">LOMv1.0</langstring>
      </source>
      <value>
        <langstring xml:lang="x-none">any</langstring>
      </value>
    </name>
    <minimumversion>0</minimumversion>
    <maximumversion>0</maximumversion>
  </requirement>
  <otherplatformrequirements>
    <langstring xml:lang="en">gLite</langstring>
  </otherplatformrequirements>
  <duration></duration>
</technical>
```

```
<educational>
  <interactivitytype>
    <source>
      <langstring xml:lang="x-none">LOMv1.0</langstring>
    </source>
    <value>
      <langstring xml:lang="x-none">active</langstring>
    </value>
  </interactivitytype>
  <learningresourcetype>
    <source>
      <langstring xml:lang="x-none">LOMv1.0</langstring>
    </source>
    <value>
      <langstring xml:lang="x-none">excercise</langstring>
    </value>
  </learningresourcetype>
  <intendedenduserrole>
    <source>
      <langstring xml:lang="x-none">LOMv1.0</langstring>
    </source>
    <value>
      <langstring xml:lang="x-none">Learner</langstring>
    </value>
  </intendedenduserrole>
  <context>
    <source>
      <langstring xml:lang="x-none">LOMv1.0</langstring>
    </source>
    <value>
      <langstring xml:lang="x-none">Higher Education</langstring>
    </value>
  </context>
  <context>
    <source>
      <langstring xml:lang="x-none">UKEC</langstring>
```

```
    </source>
    <value>
      <langstring xml:lang="x-none">Higher Education</langstring>
    </value>
  </context>
  <typicalagerange>
    <langstring xml:lang="en">18-</langstring>
  </typicalagerange>
  <difficulty>
    <source>
      <langstring xml:lang="x-none">LOMv1.0</langstring>
    </source>
    <value>
      <langstring xml:lang="x-none">medium</langstring>
    </value>
  </difficulty>
  <description>
    <langstring xml:lang="eng">
      Use to teach R-GMA DB concepts, not general DB concepts
    </langstring>
  </description>
  <language>en</language>
</educational>

<rights>
  <cost>
    <source>
      <langstring xml:lang="x-none">LOMv1.0</langstring>
    </source>
    <value>
      <langstring xml:lang="x-none">no</langstring>
    </value>
  </cost>
  <copyrightandotherrestrictions>
    <source>
      <langstring xml:lang="x-none">LOMv1.0</langstring>
    </source>
```

```
    <value>
      <langstring xml:lang="x-none">no</langstring>
    </value>
    </copyrightandotherrestrictions>
      <description>
    <langstring xml:lang="en">Freely reusable</langstring>
  </description>
</rights>

<classification>
  <purpose>
    <source>
      <langstring xml:lang="x-none">LOMv1.0</langstring>
    </source>
    <value>
      <langstring xml:lang="x-none">Discipline</langstring>
    </value>
  </purpose>
  <taxonpath>
    <source>
      <langstring xml:lang="en">LOMv1.0</langstring>
    </source>
    <taxon>
      <entry>
        <langstring xml:lang="en">Physics</langstring>
      </entry>
    </taxon>
  </taxonpath>
  <taxonpath>
    <source>
      <langstring xml:lang="en">
        http://www.cs.tcd.ie/Kathryn.Cassidy/taxonomy.html
      </langstring>
    </source>
    <taxon>
      <entry>
        <langstring xml:lang="en">Computer Science</langstring>
```

```
      </entry>
    </taxon>
  </taxonpath>
  <taxonpath>
    <source>
      <langstring xml:lang="en">
        http://www.cs.tcd.ie/Kathryn.Cassidy/taxonomy.html
      </langstring>
    </source>
    <taxon>
      <entry>
        <langstring xml:lang="en">Physics</langstring>
      </entry>
    </taxon>
  </taxonpath>
</classification>

<classification>
  <purpose>
    <source>
      <langstring xml:lang="x-none">
        http://www.cs.tcd.ie/Kathryn.Cassidy/taxonomy.html
      </langstring>
    </source>
    <value>
      <langstring xml:lang="x-none">VO</langstring>
    </value>
  </purpose>
  <taxonpath>
    <source>
      <langstring xml:lang="x-none">
        http://www.cs.tcd.ie/Kathryn.Cassidy/taxonomy.html
      </langstring>
    </source>
    <taxon>
      <entry>
        <langstring xml:lang="en">cosmo</langstring>
```

```
      </entry>
    </taxon>
  </taxonpath>
</classification>

<classification>
  <purpose>
    <source>
      <langstring xml:lang="x-none">
        http://www.cs.tcd.ie/Kathryn.Cassidy/taxonomy.html
      </langstring>
    </source>
    <value>
      <langstring xml:lang="x-none">usertype</langstring>
    </value>
  </purpose>
  <taxonpath>
    <source>
      <langstring xml:lang="x-none">
        http://www.cs.tcd.ie/Kathryn.Cassidy/taxonomy.html
      </langstring>
    </source>
    <taxon>
      <entry>
        <langstring xml:lang="en">all</langstring>
      </entry>
    </taxon>
  </taxonpath>
</classification>

</lom>
```

Appendix F

eLGrid Authentication

eLGrid uses an authenticator *Valve* for Apache Tomcat in order to allow users to log in with either their grid certificate or a username and password. Tomcat can support either one of these on its own, but cannot natively support the fallback to username and password that is required in order to allow first-time learners, who have not yet obtained a Grid certificate, to log in and complete the *Grid Certificates* course.

A Valve element in Apache Tomcat is a component that is inserted into the request processing pipeline for the associated Tomcat container. The Valve used for the eLGrid authentication is the SSLWithFormFallbackAuthenticator Valve developed by Richard Unger [225].

Using this authenticator Valve with Apache Tomcat allows the authentication to be handled entirely by the webserver without having to write any special code to handle authentication. The learner's credentials (their username or DN from their certificate) are however made available to the eLGrid application, where they can be used to identify the learner and as well as to get information about their practical exercises.

User details can be stored in a variety of formats supported by Tomcat. The simplest is a file-based system which by default stores user details in the file users.xml.

This file based solution has some drawbacks. The file is read into memory at server startup time, and if it is changed later (e.g. if a new user is added) then the server must be restarted in order for the changes to take effect. Also, it relies on the file permissions for security and this may not be enough in many scenarios.

Database authentication is an alternative that is slightly more complicated to configure but allows dynamic changes to the user information and can potentially allow more complicated permissions configuration.

SSL authentication with form fallback, authenticating via a database is configured in the Apache Tomcat server.xml file as follows (the configuration parameters are described in Table F.1):

```
<Realm className="org.apache.catalina.realm.CombinedRealm" >


  <Realm className="org.apache.catalina.realm.DataSourceRealm"
    dataSourceName="elgriddb"
    digest="MD5"
    userTable="users"
    userNameCol="userdn"
    userCredCol="userpass"
    userRoleTable="userroles"
    roleNameCol="rolename"
    localDataSource="true"/>

  <Realm className="org.apache.catalina.realm.DataSourceRealm"
    dataSourceName="elgriddb"
    digest="MD5"
    userTable="users"
    userNameCol="username"
    userCredCol="userpass"
    userRoleTable="userroles"
    roleNameCol="rolename"
    localDataSource="true"/>

</Realm>
```

The elgriddb datasource is defined in the Apache Tomcat file context.xml

```
<Resource name="elgriddb"
  auth="Container"
  type="javax.sql.DataSource"
  maxActive="100"
  maxIdle="30"
  maxWait="10000"
  username="****"
  password="****"
  driverClassName="com.mysql.jdbc.Driver"
  url="jdbc:mysql://localhost/elgrid/>

<Valve className="at.telekom.tomcat.security.SSLWithFormFallback" />
```

Table F.1: Authentication Configuration

Variable	Results
className = "org.apache.catalina.realm.CombinedRealm"	Identifies the authentication as combined authentication realm, i.e. using more than one authentication mechanism
className = "org.apache.catalina.realm.DataSourceRealm"	Identifies the individual authentication mechanisms as data source realms, that is, the credentials will be checked against a list stored in a data source specified as a resource in the file context.xml
dataSourceName = "elgriddb"	Identifies the source for the user credentials, in this case the "elgriddb" resource, which is defined later in the file
digest = "MD5"	Allows you to specify the hashing algorithm if you want to use digested passwords
userTable = "users"	The table in which user details are stored
userNameCol = "userdn" or "username"	The column of the users table in which user ids can be found, in this case it is either the userdn column or the username column, depending on whether certificate or login authentication realm was used
userCredCol = "userpass"	The column of the users table in which user passwords can be found when using login authentication(here userpass)
userRoleTable = "userroles"	The table in which the users' roles are stored (here userroles). Used to distinguish between admin and normal users.
roleNameCol = "rolename"	The column of the roles table in which the users' roles can be found (here rolename)
localDataSource ="true"	The data source is local to the Apache Tomcat machine

Continued on next page

Table F.1 – continued from previous page

Parameter	Details
auth = ”Container”	The scope of the authentication applies to a specific Apache Tomcat container
type = ”javax.sql.DataSource”	The data is stored in a SQL database
maxActive = ”100”	Maximum active connections
maxIdle = ”30”	Maximum idle connections
maxWait = ”10000”	Connection timeout value
username = ”****”	The username for connecting to the SQL database
password = ”****”	The password for connecting to the SQL database
driverClassName = ”driver”	Identifies the DB types as JDBC and tells Tomcat what driver it needs to use in order to access the database
url = ”jdbc:mysql://localhost/elgrid”	The connection url to access the database, includes the protocol, driver, host (here localhost as the Tomcat server and MySQL database reside on the same machine) and DB name (here the eLGrid database)
Valve className = ”at.telekom.tomcat.security.SSLWithFormFallback”	The name and jarfile for the SSLWithFormFallback Valve which will be used to handle this authentication

Appendix G

Data from the Second Experiment

expgroupid, learnerid, totalconcepttime, totalconceptcount, complete_initial_questionaire_count, correct_answer_count, incomplete_answer_count, incorrect_answer_count, login_count, logout_count, reset_profile_count, view_concept_count, view_test_count, elgrid-easy, future-elgrid, future-use-tools, prac-apply-concepts, prac-env-easy, prac-good-eval, practical-access, practical-understood, tests-good-eval, theory-comp-prac, too-easy, too-hard, traffic-light-prac, traffic-understood, trafficlights-nav, tutor-req, projectmark, elgridmark, totalmark, gender, level

2, 78, 10540, 89, 1, 0, 0, 2, 1, 0, 0, 96, 7, Somewhat Agree, Strongly Agree, Strongly Agree, Somewhat Agree, Somewhat Agree, No answer Given, Strongly Agree, Somewhat Agree, No answer Given, No answer Given, Strongly Disagree, Somewhat Disagree, Somewhat Disagree, Somewhat Disagree, No answer Given, Strongly Agree, NA, NA, NA, NA, NA

2, 79, 42683, 171, 2, 0, 0, 0, 2, 1, 0, 182, 7, Neither Agree nor Disagree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Strongly Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Neither Agree nor Disagree, Strongly Agree, Neither Agree nor Disagree, Neither Agree nor Disagree, Somewhat Agree, Somewhat Agree, NA, NA, NA, NA, NA

2, 81, 10450, 156, 1, 3, 0, 0, 1, 1, 0, 167, 8, Somewhat Agree, Strongly Agree, Somewhat Agree, Somewhat Agree, Neither Agree nor Disagree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Strongly Agree, Somewhat Agree, Strongly Disagree, Strongly Disagree, Somewhat Agree, Somewhat

Disagree, Somewhat Disagree, Strongly Agree, NA, NA, NA, NA, NA

2, 82, 10170, 58, 0, 0, 0, 0, 1, 0, 0, 63, 0, Somewhat Agree, Somewhat Agree, Strongly Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Disagree, Strongly Disagree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, NA, NA, NA, NA, NA

2, 83, 8895, 48, 0, 0, 0, 0, 1, 0, 0, 51, 0, Somewhat Agree, Somewhat Disagree, Somewhat Agree, Somewhat Agree, Neither Agree nor Disagree, No answer Given, Somewhat Agree, Neither Agree nor Disagree, No answer Given, Neither Agree nor Disagree, Somewhat Disagree, Strongly Disagree, No answer Given, No answer Given, No answer Given, Somewhat Agree, NA, NA, NA, NA, NA

2, 84, 10534, 42, 0, 0, 0, 0, 1, 0, 0, 47, 0, Somewhat Agree, Neither Agree nor Disagree, Somewhat Agree, Strongly Agree, Somewhat Agree, Somewhat Agree, Neither Agree nor Disagree, Strongly Agree, Neither Agree nor Disagree, Neither Agree nor Disagree, Strongly Disagree, Somewhat Agree, Neither Agree nor Disagree, Neither Agree nor Disagree, Somewhat Agree, Strongly Agree, NA, NA, NA, NA, NA

2, 85, 8774, 55, 0, 0, 0, 0, 1, 0, 0, 60, 1, Somewhat Disagree, Neither Agree nor Disagree, Strongly Disagree, Neither Agree nor Disagree, Neither Agree nor Disagree, Neither Agree nor Disagree, Somewhat Agree, Neither Agree nor Disagree, Neither Agree nor Disagree, Neither Agree nor Disagree, Neither Agree nor Disagree, Neither Agree nor Disagree, Neither Agree nor Disagree, Neither Agree nor Disagree, Neither Agree nor Disagree, Somewhat Agree, NA, NA, NA, NA, NA

Appendix H

Data from the Third Experiment

experimentid, learnerid, totalconcepttime, totalconceptcount, complete_initial_questionaire_count, correct_answer_count, incomplete_answer_count, incorrect_answer_count, login_count, logout_count, reset_profile_count, view_concept_count, view_test_count, elgrideasy, futureelgrid, futureusetools, pracapplyconcepts, pracenveasy, pracgoodeval, practicalaccess, practicalunderstood, testsgoodeval, theorycompprac, tooeasy, toohard, trafficlightprac, trafficunderstood, trafficlightsnav, tutorreq, projectmark, elgridmark, totalmark, gender, level

3, 149, 3738, 188, 2, 172, 3, 33, 8, 1, 1, 220, 22, Somewhat Agree, Strongly Disagree, Strongly Disagree, Neither, Strongly Disagree, Strongly Agree, Somewhat Disagree, Somewhat Agree, Somewhat Agree, Neither, Neither, Somewhat Disagree, Somewhat Disagree, Neither, Neither, Somewhat Agree, 90, 90, 90, M, undergrad

3, 150, 5246, 149, 2, 26, 0, 7, 9, 4, 0, 163, 8, Strongly Agree, Strongly Agree, Strongly Agree, Strongly Agree, Strongly Agree, Strongly Agree, Somewhat Agree, Strongly Agree, Strongly Agree, Strongly Agree, Strongly Disagree, Strongly Disagree, Somewhat Agree, Somewhat Agree, Strongly Agree, Neither, 70, 60, 66, M, undergrad

3, 151, 5037, 102, 2, 14, 1, 1, 10, 1, 0, 117, 6, Strongly Agree, Strongly Agree, Somewhat Agree, Strongly Agree, Somewhat Agree, Strongly Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Strongly Agree, Strongly Disagree, Strongly Disagree, Strongly Agree, Strongly Disagree, Strongly

Disagree, Strongly Disagree, 90, 100, 94, M, undergrad

3, 152, 12143, 294, 1, 7, 0, 3, 14, 12, 0, 316, 10, Strongly Agree, Somewhat Agree, Somewhat Disagree, Somewhat Agree, Neither, Somewhat Agree, Strongly Agree, Somewhat Agree, Strongly Agree, Somewhat Agree, Somewhat Disagree, Neither, Somewhat Disagree, Strongly Agree, Strongly Agree, Strongly Agree, 60, 100, 76, M, undergrad

3, 153, 92493, 198, 1, 24, 1, 2, 13, 2, 0, 218, 9, Strongly Agree, Somewhat Agree, Neither, Strongly Agree, Neither, Strongly Agree, Somewhat Agree, Strongly Agree, Strongly Agree, Strongly Agree, Strongly Disagree, Somewhat Disagree, Neither, Somewhat Agree, Somewhat Disagree, Somewhat Disagree, 90, 100, 94, M, undergrad

3, 154, 18597, 201, 3, 8, 0, 2, 11, 4, 0, 224, 13, Somewhat Agree, Neither, Somewhat Disagree, Neither, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Strongly Agree, Neither, Somewhat Agree, Strongly Disagree, Neither, Somewhat Agree, Somewhat Agree, Somewhat Agree, 70, 100, 82, M, undergrad

3, 155, 12761, 167, 3, 37, 2, 8, 13, 3, 0, 190, 19, Somewhat Agree, Somewhat Agree, Somewhat Disagree, Somewhat Agree, Neither, Neither, Somewhat Agree, Somewhat Agree, Somewhat Disagree, Somewhat Agree, Neither, Strongly Disagree, Somewhat Disagree, Somewhat Disagree, Somewhat Agree, Somewhat Agree, 70, 100, 82, M, undergrad

3, 156, 38164, 195, 3, 144, 5, 27, 11, 1, 0, 248, 29, Somewhat Agree, Somewhat Agree, Somewhat Disagree, Somewhat Agree, Strongly Agree, Somewhat Agree, Neither, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Disagree, Strongly Disagree, Somewhat Disagree, Somewhat Agree, Somewhat Agree, Strongly Disagree, 80, 90, 84, M, undergrad

3, 157, 23663, 261, 1, 19, 1, 1, 17, 2, 0, 282, 10, Strongly Agree, Somewhat Agree, Somewhat Disagree, Neither, Somewhat Disagree, Somewhat Agree, Somewhat Disagree, Strongly Agree, Neither, Strongly Agree, Strongly Agree, Somewhat Disagree, Neither, Neither, Neither, Strongly Agree, 80, 100, 88, M, undergrad

3, 158, 3307, 83, 0, 46, 3, 5, 5, 1, 0, 94, 6, Somewhat Agree, Strongly Agree, Strongly Disagree, Somewhat Agree, Neither, Strongly Agree, Somewhat Agree, Strongly Agree, Neither, Strongly Agree, Neither, Strongly Disagree, Strongly Agree, Neither, Neither, Strongly Disagree, 70, 100, 82, M, undergrad

3, 159, 12548, 103, 2, 13, 0, 7, 11, 1, 0, 120, 9, Somewhat Agree, Strongly Agree, Neither, Somewhat Agree, Somewhat Agree, Neither, Somewhat Agree, Neither, Somewhat Agree, Somewhat Agree, Neither, Strongly Disagree, Strongly Agree, Somewhat Agree, Neither, Somewhat Disagree, 80, 100, 88, M, undergrad

3, 160, 7368, 137, 1, 13, 0, 7, 8, 1, 0, 151, 9, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Neither, Somewhat Agree, Strongly Agree, Somewhat Agree, Somewhat Disagree, Strongly Disagree, Strongly Agree, Strongly Agree, Strongly Agree, Somewhat Agree, 80, 100, 88, M, undergrad

3, 161, 19072, 565, 7, 188, 3, 33, 21, 3, 5, 626, 72, Neither, Somewhat Agree, Somewhat Disagree, Somewhat Agree, Neither, Somewhat Agree, Somewhat Disagree, Neither, Somewhat Agree, Somewhat Agree, Strongly Disagree, Neither, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, 60, 90, 72, M, undergrad

3, 162, 10767, 123, 1, 19, 0, 3, 9, 3, 0, 135, 6, Somewhat Agree, Strongly Agree, Strongly Disagree, Neither, Neither, Neither, Neither, Neither, Somewhat Disagree, Somewhat Disagree, Strongly Disagree, Neither, Strongly Disagree, Somewhat Agree, Somewhat Disagree, Strongly Agree, 60, 100, 76, M, undergrad

Appendix I

Data from the Fourth & Fifth Experiments

experimentid, learnerid, totalconcepttime, totalconceptcount, complete_initial_questionaire_count, correct_answer_count, incomplete_answer_count, incorrect_answer_count, login_count, logout_count, reset_profile_count, view_concept_count, view_test_count, elgrideasy, futureelgrid, futureusetools, pracapplyconcepts, pracenveasy, pracgoodeval, practicalaccess, practicalunderstood, testsgoodeval, theorycompprac, tooeasy, toohard, trafficlightprac, trafficunderstood, trafficlightsnav, tutorreq, projectmark, elgridmark, totalmark, gender, level, xel, prac_correct_answer_count, prac_incorrect_answer_count, trafficpracunderstood, pracgoodevalsystem, trafficlightpracxel, instantfeedback

4, 164, 21343, 267, 1, 4, 0, 1, 19, 2, NA, 289, 10, Somewhat Agree, Strongly Agree, Neither, Somewhat Agree, Somewhat Agree, Somewhat Agree, Neither, Somewhat Agree, Somewhat Agree, Somewhat Agree, Neither, Neither, NA, Neither, Somewhat Agree, Strongly Agree, 100, 80, 90, M, undergrad, 1, 1, 0, Strongly Agree, Strongly Agree, Somewhat Agree, Strongly Agree

5, 181, 4998, 129, 2, 4, 0, 1, 6, 1, NA, 140, 7, NA, NA, NA, Neither, NA, NA, Neither, Somewhat Agree, NA, Somewhat Agree, Strongly Agree, Strongly Disagree, NA, NA, NA, Strongly Disagree, 90, 80, 85, M, undergrad, 1, 4, 0, NA, NA, NA, NA

5, 186, 13221, 206, 1, 9, 0, 1, 12, 1, NA, 227, 8, Somewhat Agree, Strongly Agree, Neither, Strongly Agree, Somewhat Agree, Strongly Agree, Neither, Somewhat Agree, Somewhat Agree, Neither, Somewhat Disagree, Strongly Disagree, NA, Strongly Agree, Strongly Agree, Somewhat Disagree, 100, 100, 100, M, undergrad, 1, 2, 0, Strongly Agree, Strongly Agree, Somewhat Agree, Strongly Agree

5, 185, 15419, 232, 2, 8, 0, 2, 4, 1, NA, 247, 16, Strongly Disagree, Somewhat Disagree, Strongly Disagree, Strongly Agree, Strongly Disagree, Somewhat Agree, Somewhat Disagree, Neither, Somewhat Agree, Neither, Somewhat Disagree, Somewhat Agree, NA, Neither, Strongly Disagree, Neither, 95, 100, 98, M, undergrad, 1, 2, 0, Neither, Somewhat Agree, Somewhat Disagree, Neither

5, 192, 12519, 89, 2, 5, 0, 0, 3, NA, NA, 94, 4, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, 90, 100, 95, M, undergrad, 1, 4, 0, NA, NA, NA, NA

5, 182, 23119, 387, 3, 31, 0, 3, 19, 2, NA, 422, 34, Somewhat Disagree, Neither, Strongly Disagree, Somewhat Agree, Somewhat Disagree, Strongly Agree, Somewhat Disagree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Neither, Somewhat Disagree, NA, Somewhat Agree, Somewhat Disagree, Neither, 100, 100, 100, M, undergrad, 1, 3, 0, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree

5, 191, 22077, 254, 4, 7, 0, 3, 10, 1, NA, 276, 11, Somewhat Agree, Strongly Disagree, Somewhat Disagree, Neither, Somewhat Disagree, Neither, Neither, Neither, Neither, Somewhat Disagree, Somewhat Disagree, Strongly Disagree, NA, Neither, Somewhat Agree, Strongly Disagree, 100, 80, 90, M, undergrad, 1, 2, 0, Neither, Neither, Neither, Somewhat Agree

4, 173, 11980, 118, 0, 1, 0, 0, 13, 2, 0, 131, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, 100, 0, 50, M, undergrad, 1, 1, 0, NA, NA, NA, NA

5, 180, 14102, 164, 1, 22, 0, 2, 11, 5, NA, 183, 16, Somewhat Agree, Neither, Somewhat Disagree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree, Neither, Neither, Somewhat Disagree, Strongly Disagree, NA, Strongly Agree, Somewhat Agree, Somewhat Disagree, 100, 100, 100, M, undergrad, 1, 4, 0, Strongly Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree

5, 187, 20160, 284, 3, 31, 0, 4, 9, 1, NA, 307, 19, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, 100, 100, 100, M, undergrad, 1, 6, 0, NA, NA, NA, NA

5, 188, 19396, 288, 1, 13, 0, 2, 13, 1, NA, 308, 15, Somewhat Agree, Somewhat Agree, Somewhat Agree, Strongly Agree, Neither, Somewhat Agree, Somewhat Agree, Somewhat Agree, Strongly Agree, Strongly Agree, Somewhat Disagree, Strongly Disagree, NA, Strongly Agree, Somewhat Agree, Somewhat Disagree, 90, 100, 95, M, undergrad, 1, 2, 0, Strongly Agree, Somewhat Agree, Somewhat Agree, Strongly Agree

4, 165, 7784, 161, 1, 1, 0, 0, 11, 1, 0, 174, 11, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, 100, 0, 50, M, undergrad, 1, 1, 0, NA, NA, NA, NA

4, 166, 11508, 146, 0, 5, 0, 0, 18, 10, NA, 193, 14, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, 100, 100, 100, M, undergrad, 1, 1, 0, NA, NA, NA, NA

4, 170, 24181, 236, 1, 7, 1, 2, 14, 2, NA, 254, 8, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, 100, 100, 100, M, undergrad, 1, 1, 0, NA, NA, NA, NA

4, 168, 17693, 177, 1, 5, 0, 0, 22, 2, NA, 200, 5, Somewhat Agree, Neither, Strongly Disagree, Strongly Agree, Somewhat Disagree, Somewhat Agree, Somewhat Disagree, Somewhat Agree, Somewhat Agree, Neither, Somewhat Disagree, Somewhat Disagree, NA, Strongly Agree, Somewhat Agree, Strongly Disagree, 100, 100, 100, M, undergrad, 1, 1, 0, Strongly Agree, Somewhat Agree, Neither, Strongly Agree

4, 169, 40473, 295, 1, 5, 0, 0, 17, 2, NA, 347, 4, Somewhat Agree, Strongly Agree, Neither, Somewhat Agree, Somewhat Agree, Somewhat Agree, Neither, Strongly Agree, Neither, Somewhat Disagree, Neither, Somewhat Disagree, NA, Strongly Agree, Strongly Agree, Somewhat Disagree, 100, 100, 100, M, undergrad, 1, 1, 0, Strongly Agree, Somewhat Agree, Somewhat Agree, Somewhat Agree

5, 184, 30989, 130, 1, 3, 1, 1, 17, 1, NA, 152, 4, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, 90, 70, 80, M, undergrad, 1, 2, 0, NA, NA, NA, NA

4, 171, 15547, 146, 2, 30, 0, 0, 12, 2, NA, 187, 13, Strongly Agree, Somewhat Agree, Somewhat Agree, Strongly Agree, Strongly Agree, Somewhat Agree, Strongly Agree, Strongly Agree, Somewhat Agree, Somewhat Agree, Neither, Strongly Disagree, NA, Somewhat Agree, Somewhat Agree, Somewhat Disagree, 100, 100, 100, M, undergrad, 1, 1, 0, Strongly Agree, Somewhat Agree, Somewhat Agree, Strongly Agree

4, 172, 12608, 190, 1, 5, 0, 0, 12, 3, NA, 210, 9, Somewhat Agree, Somewhat Agree, Neither, Somewhat Agree, Neither, Neither, Neither, Somewhat Agree, Somewhat Agree, Somewhat Disagree, Somewhat Disagree, Somewhat Disagree, NA, Somewhat Agree, Somewhat Agree, Strongly Disagree, 100, 100, 100, M, undergrad, 1, 2, 0, Somewhat Agree, Neither, Neither, Somewhat Agree

4, 167, 10601, 192, 3, 14, 0, 0, 9, 2, NA, 209, 16, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, 100, 100, 100, M, undergrad, 1, 1, 0, NA, NA, NA, NA

5, 183, 8716, 225, 1, 24, 1, 6, 6, 1, 2, 240, 16, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, 100, 100, 100, M, undergrad, 1, 0, 2, NA, NA, NA, NA

anon, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, Strongly Agree, Somewhat Agree, Neither, Somewhat Disagree, Strongly Agree, Strongly Agree, Strongly Agree, Strongly Agree, Strongly Agree, Neither, Neither, Strongly Disagree, NA, Somewhat Agree, Strongly Agree, Somewhat Disagree, NA, NA, NA, NA, NA, 1, NA, NA, Somewhat Agree, Strongly Agree, Neither, Somewhat Agree

Appendix J

Data from Other Learners

A number of learners, many of whom were remote learners, applied for accounts on eLGrid. These learners were analysed as group D (see chapter 5. Most made relatively little use of the system, but some were rather more active than others. End of course surveys were not administered to this group.

expgroupid, learnerid, totalconcepttime, totalconceptcount, complete_initial_questionaire_count, correct_answer_count, incomplete_answer_count, incorrect_answer_count, login_count, logout_count, reset_profile_count, view_concept_count, view_test_count, elgrid-easy, future-elgrid, future-use-tools, prac-apply-concepts, prac-env-easy, prac-good-eval, practical-access, practical-understood, tests-good-eval, theory-comp-prac, too-easy, too-hard, traffic-light-prac, traffic-understood, trafficlights-nav, tutor-req, projectmark, elgridmark, totalmark, gender, level

4, 3, 1746, 9, 0, 9, 0, 2, 3, 2, 0, 13, 1, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, postdoc

4, 8, 1246, 24, 1, 10, 0, 2, 8, 4, 0, 46, 6, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, academic

4, 11, 57, 3, 0, 0, 0, 0, 1, 0, 0, 4, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, postdoc

4, 14, 184, 13, 0, 0, 0, 0, 1, 0, 0, 14, 1, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, postdoc

4, 17, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, F, academic

4, 22, 155, 9, 0, 0, 0, 0, 1, 0, 0, 10, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, postdoc

4, 24, 3636, 35, 0, 0, 0, 0, 5, 0, 0, 42, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, postgrad

4, 29, 3254, 105, 0, 0, 0, 0, 2, 1, 0, 111, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, postdoc

4, 33, 214, 5, 0, 0, 0, 0, 2, 1, 0, 7, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, NA

4, 34, 72530, 194, 1, 4, 0, 2, 10, 7, 0, 224, 3, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, postgrad

4, 65, 2625, 38, 0, 0, 0, 0, 1, 0, 0, 44, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA

4, 67, 3669, 45, 0, 0, 0, 0, 1, 0, 0, 53, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA

4, 68, 6516, 68, 0, 0, 0, 0, 6, 1, 1, 85, 1, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA

4, 69, 759, 36, 0, 0, 0, 0, 2, 1, 0, 52, 1, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA

4, 70, 6915, 45, 0, 0, 0, 0, 4, 1, 0, 52, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA

4, 71, 8735, 116, 0, 0, 0, 0, 10, 2, 0, 148, 1, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA

4, 72, 3397, 46, 1, 11, 0, 1, 4, 0, 0, 63, 2, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA

4, 73, 2629, 39, 0, 0, 0, 0, 3, 2, 0, 44, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA

4, 74, 8257, 76, 0, 0, 0, 0, 2, 0, 0, 87, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA

4, 76, 293029, 377, 1, 0, 0, 0, 28, 0, 0, 457, 3, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, NA

4, 80, 2256, 34, 0, 0, 0, 0, 3, 2, 1, 39, 2, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, postgrad

4, 102, 275, 9, 0, 0, 0, 0, 1, 0, 0, 10, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA

4, 110, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, NA

4, 111, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, NA

4, 112, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, NA

4, 113, 0, 0, 0, 0, 0, 0, 2, 2, 0, 0, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA

4, 114, 3553, 108, 1, 8, 0, 3, 3, 1, 0, 126, 3, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, NA

4, 115, 0, 0, 0, 0, 0, 0, 2, 1, 0, 1, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, academic

4, 116, 18, 3, 0, 0, 0, 0, 2, 1, 0, 5, 1, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, NA

4, 117, 0, 0, 0, 0, 0, 0, 1, 1, 0, 15, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, academic

4, 118, 3013, 3, 0, 0, 0, 0, 1, 1, 0, 6, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA

4, 124, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, postgrad

4, 133, 5104, 67, 0, 0, 0, 0, 4, 3, 0, 73, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, postgrad

4, 135, 14, 1, 0, 0, 0, 0, 1, 1, 0, 4, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, academic

4, 136, 48723, 176, 1, 11, 0, 0, 10, 1, 0, 204, 2, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, postgrad

4, 137, 350, 28, 2, 11, 0, 0, 3, 2, 0, 33, 3, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, postdoc

4, 138, 214, 34, 0, 0, 0, 0, 2, 1, 0, 36, 1, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, M, NA

4, 163, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, F, academic

Appendix K

Treatment of Missing Data

Because of the small sample size, where possible, for example when analysing responses to individual survey questions, pair-wise deletion has been used. This means that only the missing item has been discarded, not all responses from that individual.

For correlations listwise deletion has been employed and all responses from learners with missing data have been discarded.

For the data mining analysis, missing data was treated as a separate category.

Where the missing data was not considered to be indicative of a learner attitude it was treated as if the learner had not been asked that question, for example, some learners only filled in one page of the feedback form, it was assumed that their missing responses to the subsequent pages could be treated as if they had not seen the questions.

A list of unanswered questions is given in table K.1.

Table K.1: Missing survey data points

Learner	Statement
78	The theory in this course supported and complemented the practical exercises
78	The traffic-light indicators helped me to navigate the courses
78	The tests were a good way for the system to evaluate my knowledge
78	The practical exercises were a good way for me to evaluate my knowledge
83	The traffic-light indicators helped me to navigate the courses

Continued on next page

Table K.1 – continued from previous page

Learner	Statement
83	I understood how my test answers affected the traffic-light indicators
83	The tests were a good way for the system to evaluate my knowledge
83	The practical exercises were a good way for me to evaluate my knowledge
83	It would be useful if the traffic-light indicators took into account the results of my practical exercises
75	I understood how my test answers affected the traffic-light indicators
75	The tests were a good way for the system to evaluate my knowledge
9	The practical exercises helped me to understand and apply the course concepts
9	The theory in this course supported and complemented the practical exercises
9	The presence of a tutor was necessary for me to complete this course
61	Accessing the practical environments from within the eLGrid eLearning system made the practicals easier to complete
181	The eLGrid eLearning system is easy to use
181	The practical environment (WebCom IDE) is easy to use
181	The ”traffic-light” indicators helped me to navigate the course
181	I understood how my test answers affected the traffic-light indicators
181	I understood how the results of my practical exercises affected the traffic-light indicators
181	The tests were a good way for the system to evaluate my knowledge
181	The practical exercises were a good way for me to evaluate my knowledge
181	The practical exercises which collected my results were a good way for the system to evaluate my knowledge
181	The fact that the traffic-light indicators took into account the results of the practical exercises as well as my test answers was helpful to me
181	The instant feedback on the results of the final practical exercise was helpful to me
181	In the future I will probably use the technologies taught in this course (WebCom)
181	If I need to learn about Grid technologies in future I would use the eLGrid eLearning system

Appendix L

Tables of Statistical Results

Table L.1: ANOVA results

<table>
<tr><th>Variable</th><th>Results</th></tr>
<tr><td>Concept count by groups</td><td>
ANOVA
<table>
<tr><th></th><th>Df</th><th>Sum Sq</th><th>Mean Sq</th><th>F value</th><th>Pr(>F)</th></tr>
<tr><td>Group ID</td><td>3</td><td>236393</td><td>78798</td><td>12.183</td><td>1.667e-06</td></tr>
<tr><td>Residuals</td><td>70</td><td>452744</td><td>6468</td><td></td><td></td></tr>
</table>
Significant difference found
Multiple Comparisons
alpha = 0.05, df = 70
<table>
<tr><th></th><th>statistic</th><th>critical value</th><th>Significant</th></tr>
<tr><td>A vs. B</td><td>0.31</td><td>2.7153</td><td>No</td></tr>
<tr><td>A vs. C</td><td>4.048729</td><td>2.7153</td><td>Yes</td></tr>
<tr><td>A vs. D</td><td>1.268185</td><td>2.7153</td><td>No</td></tr>
<tr><td>B vs. C</td><td>2.943992</td><td>2.7153</td><td>Yes</td></tr>
<tr><td>B vs. D</td><td>1.284832</td><td>2.7153</td><td>No</td></tr>
<tr><td>C vs. D</td><td>6.049313</td><td>2.7153</td><td>Yes</td></tr>
</table>
</td></tr>
</table>

Continued on next page

Table L.1 – continued from previous page

<table>
<tr><th>Variable</th><th>Results</th></tr>
<tr><td>Login count by groups</td><td>
ANOVA
<table>
<tr><th></th><th>Df</th><th>Sum Sq</th><th>Mean Sq</th><th>F value</th><th>Pr(>F)</th></tr>
<tr><td>expgroupid</td><td>3</td><td>766.51</td><td>255.505</td><td>14.262</td><td>2.387e-07</td></tr>
<tr><td>Residuals</td><td>70</td><td>1254.08</td><td>17.915</td><td></td><td></td></tr>
</table>
Significant difference found
Multiple Comparisons
alpha = 0.01, df = 70
<table>
<tr><th></th><th>statistic</th><th>critical value</th><th>Significant</th></tr>
<tr><td>A vs. B</td><td>2.405224</td><td>3.2707</td><td>No</td></tr>
<tr><td>A vs. C</td><td>3.560287</td><td>3.2707</td><td>Yes</td></tr>
<tr><td>A vs. D</td><td>1.743285</td><td>3.2707</td><td>No</td></tr>
<tr><td>B vs. C</td><td>5.236451</td><td>3.2707</td><td>Yes</td></tr>
<tr><td>B vs. D</td><td>1.384325</td><td>3.2707</td><td>No</td></tr>
<tr><td>C vs. D</td><td>5.932122</td><td>3.2707</td><td>Yes</td></tr>
</table>
</td></tr>
</table>

Continued on next page

Table L.1 – continued from previous page

<table>
<tr><th>Variable</th><th>Results</th></tr>
<tr><td>Tutor required by groups</td><td>
ANOVA
<table>
<tr><th></th><th>Df</th><th>Sum Sq</th><th>Mean Sq</th><th>F value</th><th>Pr(>F)</th></tr>
<tr><td>expgroupid</td><td>2</td><td>15.847</td><td>7.923</td><td>4.6877</td><td>0.01826</td></tr>
<tr><td>Residuals</td><td>26</td><td>43.946</td><td>1.690</td><td></td><td></td></tr>
</table>
Significant difference found
Multiple Comparisons
alpha = 0.05, df = 26
<table>
<tr><th></th><th>statistic</th><th>critical value</th><th>Significant</th></tr>
<tr><td>A vs. B</td><td>3.061544</td><td>2.5589</td><td>Yes</td></tr>
<tr><td>A vs. C</td><td>1.457813</td><td>2.5589</td><td>No</td></tr>
<tr><td>B vs. C</td><td>2.027160</td><td>2.5589</td><td>No</td></tr>
</table>
</td></tr>
</table>

Continued on next page

Table L.1 – continued from previous page

<table>
<tr><th>Variable</th><th>Results</th></tr>
<tr><td>View Test Count by groups</td><td>
ANOVA
<table>
<tr><th></th><th>Df</th><th>Sum Sq</th><th>Mean Sq</th><th>F value</th><th>Pr(>F)</th></tr>
<tr><td>expgroupid</td><td>3</td><td>2602.7</td><td>867.6</td><td>14.516</td><td>1.894e-07</td></tr>
<tr><td>Residuals</td><td>70</td><td>4183.6</td><td>59.8</td><td></td><td></td></tr>
</table>
Significant difference found
Multiple Comparisons
alpha = 0.01, df = 70
<table>
<tr><th></th><th>statistic</th><th>critical value</th><th>Significant</th></tr>
<tr><td>A vs. B</td><td>0.5340959</td><td>3.2707</td><td>No</td></tr>
<tr><td>A vs. C</td><td>5.18646</td><td>3.2707</td><td>Yes</td></tr>
<tr><td>A vs. D</td><td>0.2477345</td><td>3.2707</td><td>No</td></tr>
<tr><td>B vs. C</td><td>3.635422</td><td>3.2707</td><td>Yes</td></tr>
<tr><td>B vs. D</td><td>0.7780516</td><td>3.2707</td><td>No</td></tr>
<tr><td>C vs. D</td><td>6.406362</td><td>3.2707</td><td>Yes</td></tr>
</table>
</td></tr>
</table>

Continued on next page

Table L.1 – continued from previous page

<table>
<tr><th>Variable</th><th>Results</th></tr>
<tr><td>Correct Answer Count by groups</td><td>
ANOVA
<table>
<tr><th></th><th>Df</th><th>Sum Sq</th><th>Mean Sq</th><th>F value</th><th>Pr(>F)</th></tr>
<tr><td>expgroupid</td><td>3</td><td>28796</td><td>9599</td><td>12.222</td><td>1.607e-06</td></tr>
<tr><td>Residuals</td><td>70</td><td>54977</td><td>785</td><td></td><td></td></tr>
</table>
Significant difference found
Multiple Comparisons
alpha = 0.01, df = 70
<table>
<tr><th></th><th>statistic</th><th>critical value</th><th>Significant</th></tr>
<tr><td>A vs. B</td><td>0.1793743</td><td>3.2707</td><td>No</td></tr>
<tr><td>A vs. C</td><td>4.740608</td><td>3.2707</td><td>Yes</td></tr>
<tr><td>A vs. D</td><td>0.1228709</td><td>3.2707</td><td>No</td></tr>
<tr><td>B vs. C</td><td>3.983002</td><td>3.2707</td><td>Yes</td></tr>
<tr><td>B vs. D</td><td>0.1085306</td><td>3.2707</td><td>No</td></tr>
<tr><td>C vs. D</td><td>5.754628</td><td>3.2707</td><td>Yes</td></tr>
</table>
</td></tr>
</table>

Continued on next page

Table L.1 – continued from previous page

Variable	Results
Incorrect Answer Count by groups	(see below)

ANOVA

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
expgroupid	3	1036.69	345.56	12.687	1.031e-06
Residuals	70	1906.66	27.24		

Significant difference found

Multiple Comparisons

alpha = 0.05, df = 70

	statistic	critical value	Significant
B vs. B	0.03484585	2.7153	No
A vs. C	0.8831147	2.7153	No
A vs. D	0.05499936	2.7153	No
B vs. C	0.7433971	2.7153	No
B vs. D	0.001995366	2.7153	No
C vs. D	1.103332	2.7153	No

Continued on next page

Table L.1 – continued from previous page

Variable	Results					
Incorrect Answer Count by level	**ANOVA**					
		Df	Sum Sq	Mean Sq	F value	Pr(>F)
	level	4	1034.5	258.614	9.348	4.266e-06
	Residuals	69	1908.9	27.665		
	Significant difference found					
	Multiple Comparisons					
	alpha = 0.5, df = 30					
	Academiv vs. Postdoc	-0.01090286		2.8247		No
	Academic vs. Postgrad	0.01134805		2.8247		No
	Academic vs. Undergrad	-2.688747		2.8247		No
	Postdoc vs. Postgrad	0.01090286		2.8247		No
	Postdoc vs. Undergrad	-2.538339		2.8247		No
	Postgrad vs. Undergrad	-2.688747		2.8247		No

Continued on next page

Table L.1 – continued from previous page

<table>
<tr><th>Variable</th><th>Results</th></tr>
<tr><td>Incomplete Answer count by groups</td><td>
ANOVA
<table>
<tr><th></th><th>Df</th><th>Sum Sq</th><th>Mean Sq</th><th>F value</th><th>Pr(>F)</th></tr>
<tr><td>expgroupid</td><td>3</td><td>20.907</td><td>6.969</td><td>14.688</td><td>1.623e-07</td></tr>
<tr><td>Residuals</td><td>70</td><td>33.214</td><td>0.474</td><td></td><td></td></tr>
</table>
Significant difference found
Multiple Comparisons
alpha = 0.05, df = 70
<table>
<tr><th></th><th>statistic</th><th>critical value</th><th>Significant</th></tr>
<tr><td>A vs. B</td><td>1.782829e-16</td><td>2.7153</td><td>No</td></tr>
<tr><td>A vs. C</td><td>0.1305900</td><td>2.7153</td><td>No</td></tr>
<tr><td>A vs. D</td><td>1.310623e-16</td><td>2.7153</td><td>No</td></tr>
<tr><td>B vs. C</td><td>0.1048341</td><td>2.7153</td><td>No</td></tr>
<tr><td>B vs. D</td><td>1.012432e-16</td><td>2.7153</td><td>No</td></tr>
<tr><td>C vs. D</td><td>0.1552220</td><td>2.7153</td><td>No</td></tr>
</table>
</td></tr>
</table>

Continued on next page

Table L.1 – continued from previous page

<table>
<tr><th>Variable</th><th>Results</th></tr>
<tr><td>Incomplete Answer count by level</td><td>
ANOVA
<table>
<tr><th></th><th>Df</th><th>Sum Sq</th><th>Mean Sq</th><th>F value</th><th>Pr(>F)</th></tr>
<tr><td>level</td><td>4</td><td>20.907</td><td>5.2268</td><td>10.858</td><td>6.928e-07</td></tr>
<tr><td>Residuals</td><td>69</td><td>33.214</td><td>0.4814</td><td></td><td></td></tr>
</table>
Significant difference found
Multiple Comparisons
alpha = 0.05, df = 30
<table>
<tr><td>Academiv vs. Postdoc</td><td>-7.857941e-16</td><td>2.8247</td><td>No</td></tr>
<tr><td>Academic vs. Postgrad</td><td>-5.351783e-16</td><td>2.8247</td><td>No</td></tr>
<tr><td>Academic vs. Undergrad</td><td>2.79216</td><td>2.8247</td><td>No</td></tr>
<tr><td>Postdoc vs. Postgrad</td><td>2.716115e-16</td><td>2.8247</td><td>No</td></tr>
<tr><td>Postdoc vs. Undergrad</td><td>-2.648875</td><td>2.8247</td><td>No</td></tr>
<tr><td>Postgrad vs. Undergrad</td><td>-2.79216</td><td>2.8247</td><td>No</td></tr>
</table>
</td></tr>
</table>

Continued on next page

Table L.1 – continued from previous page

<table>
<tr><th>Variable</th><th>Results</th></tr>
<tr><td>Practical Access by groups</td><td>
ANOVA
<table>
<tr><th></th><th>Df</th><th>Sum Sq</th><th>Mean Sq</th><th>F value</th><th>Pr(>F)</th></tr>
<tr><td>expgroupid</td><td>2</td><td>4.8344</td><td>2.4172</td><td>3.632</td><td>0.04064</td></tr>
<tr><td>Residuals</td><td>26</td><td>17.3036</td><td>0.6655</td><td></td><td></td></tr>
</table>
Significant difference found
Multiple Comparisons
alpha = 0.05, df = 70
<table>
<tr><th></th><th>statistic</th><th>critical value</th><th>Significant</th></tr>
<tr><td>A vs. B</td><td>0.8763325</td><td>2.5589</td><td>No</td></tr>
<tr><td>A vs. C</td><td>2.599826</td><td>2.5589</td><td>Yes</td></tr>
<tr><td>B vs. C</td><td>1.509375</td><td>2.5589</td><td>No</td></tr>
</table>
</td></tr>
<tr><td>"In the future I will probably use the technologies taught in this course" by groups</td><td>
t-test
<table>
<tr><th>t</th><th>degrees of freedom</th><th>p-value</th><th>95% Confidence Interval</th></tr>
<tr><td>2.3638</td><td>11.105</td><td>0.03736</td><td>[0.09990671, 2.75723614]</td></tr>
</table>
Significant difference found at alpha = 0.05
</td></tr>
</table>

Continued on next page

Table L.1 – continued from previous page

<table>
<tr><th>Variable</th><th>Results</th></tr>
<tr><td>Total concept count by level with unknowns removed</td><td>

ANOVA

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
level	3	194510	64837	7.3633	0.0007723
Residuals	30	264161	8805		

Significant difference found

Multiple Comparisons

alpha = 0.01 df = 30

	statistic	critical value	Significant
Academic vs. Postdoc	0.04788722	3.4543	No
Academiv vs. Postgrad	1.232109	3.4543	No
Academic vs. Undergrad	3.860674	3.4543	Yes
Postdoc vs Postgrad	1.231659	3.4543	No
Postdoc vs. Undergrad	3.717157	3.4543	Yes
Postgrad vs. Undergratuate	2.437957	3.4543	No

</td></tr>
<tr><td>Total concept time by groups</td><td>

ANOVA

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
expgroupid	3	1.2969e+09	432310554	0.3157	0.814
Residuals	70	9.5866e+10	1369519191		

No significant difference found

</td></tr>
</table>

Continued on next page

Table L.1 – continued from previous page

<table>
<tr><th>Variable</th><th>Results</th></tr>
<tr><td>Total concept time by level</td><td>
ANOVA
<table>
<tr><th></th><th>Df</th><th>Sum Sq</th><th>Mean Sq</th><th>F value</th><th>Pr(>F)</th></tr>
<tr><td>level</td><td>4</td><td>2.6422e+09</td><td>660561992</td><td>0.4822</td><td>0.7487</td></tr>
<tr><td>Residuals</td><td>69</td><td>9.4521e+10</td><td>1369869958</td><td></td><td></td></tr>
</table>
No significant difference found
</td></tr>
<tr><td>Logs of Total Concept Time by groups (zero times removed)</td><td>
ANOVA
<table>
<tr><th></th><th>Df</th><th>Sum Sq</th><th>Mean Sq</th><th>F value</th><th>Pr(>F)</th></tr>
<tr><td>expgroupid</td><td>3</td><td>49.866</td><td>16.622</td><td>5.9398</td><td>0.001289</td></tr>
<tr><td>Residuals</td><td>60</td><td>167.905</td><td>2.798</td><td></td><td></td></tr>
</table>
Significant difference found
Multiple Comparisons
alpha = 0.05, df=60
<table>
<tr><th></th><th>statistic</th><th>critical value</th><th>Significant</th></tr>
<tr><td>A vs. B</td><td>1.213879</td><td>2.7285</td><td>No</td></tr>
<tr><td>A vs. C</td><td>1.486693</td><td>2.7285</td><td>No</td></tr>
<tr><td>A vs. D</td><td>1.965427</td><td>2.7285</td><td>No</td></tr>
<tr><td>B vs. C</td><td>0</td><td>2.7285</td><td>No</td></tr>
<tr><td>B vs. D</td><td>2.853232</td><td>2.7285</td><td>Yes</td></tr>
<tr><td>C vs. D</td><td>3.692063</td><td>2.7285</td><td>Yes</td></tr>
</table>
</td></tr>
</table>

Continued on next page

Table L.1 – continued from previous page

Variable	Results
Logs of Total Concept Time by level (zero times removed)	see below

ANOVA

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
level	3	67.810	22.603	10.389	0.0001271
Residuals	25	54.396	2.176		

Significant difference found

Multiple Comparisons

alpha = 0.05, df = 25

	statistic	critical value	Significant
Academic vs. Postdoc	0.2780361	2.8649	No
Academiv vs. Postgrad	2.847474	2.8649	No
Academic vs. Undergrad	3.367251	2.8649	Yes
Postdoc vs Postgrad	3.827952	2.8649	Yes
Postdoc vs. Undergrad	4.793271	2.8649	Yes
Postgrad vs. Undergratuate	0.2639772	2.8649	No

Continued on next page

Table L.1 – continued from previous page

<table>
<tr><th>Variable</th><th>Results</th></tr>
<tr><td>Login count by level (where level is known)</td><td>

ANOVA

<table>
<tr><th></th><th>Df</th><th>Sum Sq</th><th>Mean Sq</th><th>F value</th><th>Pr(>F)</th></tr>
<tr><td>level</td><td>3</td><td>543.91</td><td>181.30</td><td>14.622</td><td>4.755e-06</td></tr>
<tr><td>Residuals</td><td>30</td><td>371.98</td><td>12.40</td><td></td><td></td></tr>
</table>

Significant difference found

Multiple Comparisons

alpha=0.05, df=30

<table>
<tr><th></th><th>statistic</th><th>critical value</th><th>Significant</th></tr>
<tr><td>Academic vs. Postdoc</td><td>0.6686938</td><td>2.8247</td><td>No</td></tr>
<tr><td>Academic vs. Postgrad</td><td>1.593890</td><td>2.8247</td><td>No</td></tr>
<tr><td>Academic vs. Undergrad</td><td>5.067415</td><td>2.8247</td><td>Yes</td></tr>
<tr><td>Postdoc vs. Postgrad</td><td>-2.200054</td><td>2.8247</td><td>No</td></tr>
<tr><td>Postdoc vs. Undergrad</td><td>5.5698</td><td>2.8247</td><td>Yes</td></tr>
<tr><td>Postgrad vs. Undergrad</td><td>3.226949</td><td>2.8247</td><td>Yes</td></tr>
</table>

</td></tr>
</table>

Continued on next page

Table L.1 – continued from previous page

<table>
<tr><th>Variable</th><th>Results</th></tr>
<tr><td>Total concept count by level (where level is known)</td><td>
ANOVA
<table>
<tr><th></th><th>Df</th><th>Sum Sq</th><th>Mean Sq</th><th>F value</th><th>Pr(>F)</th></tr>
<tr><td>level</td><td>3</td><td>194510</td><td>64837</td><td>7.3633</td><td>0.0007723</td></tr>
<tr><td>Residuals</td><td>30</td><td>264161</td><td>8805</td><td></td><td></td></tr>
</table>
Multiple Comparisons
alpha=0.01, df=30
<table>
<tr><th></th><th>statistic</th><th>critical value</th><th>Significant</th></tr>
<tr><td>Academic vs. Postdoc</td><td>0.04788722</td><td>3.4543</td><td>No</td></tr>
<tr><td>Academic vs. Postgrad</td><td>1.232109</td><td>3.4543</td><td>No</td></tr>
<tr><td>Academic vs. Undergrad</td><td>3.860674</td><td>3.4543</td><td>Yes</td></tr>
<tr><td>Postdoc vs. Postgrad</td><td>1.231659</td><td>3.4543</td><td>No</td></tr>
<tr><td>Postdoc vs. Undergrad</td><td>3.717157</td><td>3.4543</td><td>Yes</td></tr>
<tr><td>Postgrad vs. Undergrad</td><td>2.437957</td><td>3.4543</td><td>No</td></tr>
</table>
</td></tr>
</table>

Continued on next page

Table L.1 – continued from previous page

<table>
<tr><th>Variable</th><th>Results</th></tr>
<tr><td>View Test Count by level (where level is known)</td><td>

ANOVA

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
level	3	1925.1	641.7	4.8272	0.007369
Residuals	30	3988.0	132.9		

Significant difference found

Multiple Comparisons

alpha=0.05, df=30

	statistic	critical value	Significant
Academic vs. Postdoc	0.003741523	2.8247	No
Academic vs. Postgrad	-0.07025968	2.8247	No
Academic vs. Undergrad	-2.893472	2.8247	Yes
Postdoc vs. Postgrad	-0.07124484	2.8247	No
Postdoc vs. Undergrad	-2.749255	2.8247	No
Postgrad vs. Undergrad	-2.812343	2.8247	No

</td></tr>
</table>

Continued on next page

Variable	Results
Correct Answer Count by level (where level is known)	See below

Correct Answer Count by level (where level is known)

ANOVA

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
level	3	20389	6796	3.7739	0.02073
Residuals	30	54026	1801		

Significant difference found

Multiple Comparisons

alpha=0.05, df=30

	statistic	critical value	Significant
Academic vs. Postdoc	-0.08047553	2.8247	No
Academic vs. Postgrad	-0.05025695	2.8247	No
Academic vs. Undergrad	-2.579358	2.8247	No
Postdoc vs. Postgrad	0.03219021	2.8247	No
Postdoc vs. Undergrad	-2.355237	2.8247	No
Postgrad vs. Undergrad	-2.521326	2.8247	No

Responses to the statement "The tests were a good way for the system to evaluate my knowledge" by groups

ANOVA

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
expgroupid	2	1.4981	0.74905	0.7146	0.5004
Residuals	22	23.0619	1.04827		

No significant difference found

Continued on next page

Table L.1 – continued from previous page

Variable: The theory in this course supported and complemented the practical exercises

Results:

ANOVA

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
expgroupid	2	1.381	0.69048	0.8593	0.4361
Residuals	24	19.286	0.80357		

No significant difference found

Variable: The practical exercises helped me to understand and apply the course concepts

Results:

ANOVA

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
expgroupid	2	1.2143	0.60714	1.4262	0.2591
Residuals	25	10.6429	0.42571		

No significant difference found

Variable: The practical instructions were easy to understand

Results:

ANOVA

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
expgroupid	2	1.0289	0.51447	0.7574	0.4790
Residuals	26	17.6607	0.67926		

No significant difference found

Variable: The practical environment was easy to use

Results:

ANOVA

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
expgroupid	2	0.3214	0.16071	0.1508	0.8608
Residuals	25	26.6429	1.06571		

No significant difference found

Continued on next page

Table L.1 – continued from previous page

Variable	Results					
The practical exercises were a good way for me to evaluate my knowledge	**ANOVA**					
		Df	Sum Sq	Mean Sq	F value	Pr(>F)
	expgroupid	2	0.0824	0.04121	0.0754	0.9276
	Residuals	23	12.5714	0.54658		
	No significant difference found					
The course expected too much prior knowledge	**ANOVA**					
		Df	Sum Sq	Mean Sq	F value	Pr(>F)
	expgroupid	2	3.830	1.9150	1.849	0.1775
	Residuals	26	26.929	1.0357		
	No significant difference found					
The course was too simple, explaining things that I already knew	**ANOVA**					
		Df	Sum Sq	Mean Sq	F value	Pr(>F)
	expgroupid	2	1.2728	0.63639	0.5788	0.5677
	Residuals	26	28.5893	1.09959		
	No significant difference found					
The eLGrid system is easy to use	**ANOVA**					
		Df	Sum Sq	Mean Sq	F value	Pr(>F)
	expgroupid	2	2.9643	1.4821	2.3158	0.1195
	Residuals	25	16.0000	0.6400		
	No significant difference found					

Continued on next page

Table L.1 – continued from previous page

Variable	Results
The traffic-light indicators helped me to navigate the course	**ANOVA**

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
expgroupid	2	3.696	1.8478	1.2375	0.3087
Residuals	23	34.343	1.4932		

No significant difference found

Variable	Results
I understood how my test answers affected the traffic-light indicators	**ANOVA**

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
expgroupid	2	4.4432	2.2216	1.9964	0.1587
Residuals	23	25.5952	1.1128		

No significant difference found

Variable	Results
It would be useful if the traffic-light indicators took into account the results of my practical exercises as well as my	**ANOVA**

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
expgroupid	1	0.0595	0.05952	0.0387	0.8463
Residuals	18	27.6905	1.53836		

No significant difference found

Variable	Results
If I need to learn about Grid technologies in future I would use the eLGrid eLearning system	**t-test**

t	degrees of freedom	p-value	95% Confidence Interval
-0.7019	11.695	0.4965	[-1.4689811, 0.7546954]

No significant difference found

Continued on next page

Table L.1 – continued from previous page

<table>
<tr><th>Variable</th><th>Results</th></tr>
<tr><td>"It would be useful if the traffic-light indicators took into account the results of my practical exercises as well as my test answers" By response to "The traffic-light indicators helped me to navigate the course"</td><td>ANOVA
<table>
<tr><th></th><th>Df</th><th>Sum Sq</th><th>Mean Sq</th><th>F value</th><th>Pr(>F)</th></tr>
<tr><td>trafficlights.nav</td><td>1</td><td>0.1029</td><td>0.10294</td><td>0.0673</td><td>0.7984</td></tr>
<tr><td>Residuals</td><td>17</td><td>26.0023</td><td>1.52955</td><td></td><td></td></tr>
</table>
No significant difference found</td></tr>
</table>

Table L.2: Correlations

Variables	**Pearson's correlation coefficient**
Time spent viewing concepts and total concept count	0.5300384
Login count and total concept count	0.8486979
Login count and time spent viewing concepts	0.6393305
Correct answer count and total concept count	0.574258
Total concept count and view test count	0.7440982
Correct answer count and incorrect answer count	0.9782124
Correct answer count and incorrect answer count as ratios of tests viewed	0.2062116
Responses to the statement "The tests were a good way for the system to evaluate my knowledge" and View test count	0.4522491
Responses to the statement "The tests were a good way for the system to evaluate my knowledge" and Correct Answer Count	0.380614
Responses to the statement "The tests were a good way for the system to evaluate my knowledge" and Incorrect Answer Count	0.3934024
Responses to the statement "The tests were a good way for the system to evaluate my knowledge" and Incomplete Answer Count	0.3792375
Time spent viewing concepts and total mark (group C only)	0.3477104
Concept count and total mark (group C only)	-0.3696916

Table L.3: Results for the analysis of the XeL and non-XeL group using the WebCom course

Variable	Results
Login Counts for XeL and non-XeL groups	**t-test** t / degrees of freedom / p-value / 95% Confidence Interval 0.5189 / 31.967 / 0.6074 / [-2.368570 3.987617] No significant difference found
Login Durations for XeL and non-XeL groups	**t-test** t / degrees of freedom / p-value / 95% Confidence Interval 1.3989 / 163.182 / 0.1637 / [-46369.93 271707.08] No significant difference found
Time spent viewing concepts for XeL and non-XeL groups	**t-test** t / degrees of freedom / p-value / 95% Confidence Interval -0.2873 / 15.211 / 0.7777 / [-15585.30 11878.44] No significant difference found
Number of views of the ”Recursion” topic for XeL and non-XeL groups	**t-test** t / degrees of freedom / p-value / 95% Confidence Interval 2.8702 / 29.501 / 0.007514 / [2.10485 12.51420] Significant difference found at alpha = 0.01
View Test Count for XeL and non-XeL groups	**t-test** t / degrees of freedom / p-value / 95% Confidence Interval -0.9878 / 16.069 / 0.3379 / [-15.277688 5.563402] No significant difference found

Continued on next page

Table L.3 – continued from previous page

Variable	Results
Correct Answer Count for XeL and non-XeL groups	**t-test** t / degrees of freedom / p-value / 95% Confidence Interval 0.5804 / 26.247 / 0.5666 / [-5.442322 9.728036] No significant difference found
Incorrect Answer Count for XeL and non-XeL groups	**t-test** t / degrees of freedom / p-value / 95% Confidence Interval -2.7227 / 13.332 / 0.0171 / [-1.792517 -15.397959] Significant difference found at alpha = 0.05
Total Marks for XeL and non-XeL groups	**t-test** t / degrees of freedom / p-value / 95% Confidence Interval 2.2863 / 32.011 / 0.02900 / [0.987088 17.108150] Significant difference found at alpha = 0.05
Practical Marks for XeL and non-XeL groups	**t-test** t / degrees of freedom / p-value / 95% Confidence Interval 7.4956 / 15.413 / 1.609e-06 / [16.37267 29.34162] Significant difference found at alpha = 0.01
eLGrid Marks for XeL and non-XeL groups	**t-test** t / degrees of freedom / p-value / 95% Confidence Interval -1.2264 / 27.062 / 0.2306 / [-23.546952 5.927905] No significant difference found

Continued on next page

Table L.3 – continued from previous page

Variable	Results
Responses to the statement ”The presence of a tutor was required in order for me to complete this course” for XeL and non-XeL groups	**t-test** t: -2.2201 · degrees of freedom: 23.731 · p-value: 0.03623 · 95% Confidence Interval: [0.07937929 2.19534599] Significant difference found at alpha = 0.05
Responses to the statement ”The course expected too much prior knowledge” for XeL and non-XeL groups	**t-test** t: 0.143 · degrees of freedom: 24.094 · p-value: 0.8875 · 95% Confidence Interval: [-0.6643491 0.7632502] No significant difference found
Responses to the statement ”The course was too simple, explaining things that I already knew” for XeL and non-XeL groups	**t-test** t: 0.7931 · degrees of freedom: 23.101 · p-value: 0.4358 · 95% Confidence Interval: [-0.5300098 1.1893505] No significant difference found
Responses to the statement ”The practical instructions were easy to understand” for XeL and non-XeL groups	**t-test** t: 0.0208 · degrees of freedom: 24.929 · p-value: 0.9836 · 95% Confidence Interval: [-0.5380345 0.5490235] No significant difference found

Continued on next page

Table L.3 – continued from previous page

Variable	**Results**			
Responses to the statement ”The practical environment was easy to access” for XeL and non-XeL groups	**t-test** t -0.5255 No significant difference found	degrees of freedom 24.423	p-value 0.604	95% Confidence Interval [-0.9739781 0.5783737]
Responses to the statement ”The practical exercises helped me to understand and apply the course concepts” for XeL and non-XeL groups	**t-test** t 0.4512 No significant difference found	degrees of freedom 22.466	p-value 0.6561	95% Confidence Interval [-0.5326541 0.8293574]
Responses to the statement ”The eL-Grid eLearning application is easy to use” for XeL and non-XeL groups	**t-test** t -1.46 No significant difference found	degrees of freedom 16.27	p-value 0.1633	95% Confidence Interval [-1.3125380 0.2411094]
Responses to the statement ”The traffic-light indicators helped me to navigate the course” for XeL and non-XeL groups	**t-test** t 0.8525 No significant difference found	degrees of freedom 23.559	p-value 0.4025	95% Confidence Interval [-0.5761527 1.3856765]

Continued on next page

Table L.3 – continued from previous page

Variable	Results
Responses to the statement ”In the future I will probably use the technologies taught by this course” for XeL and non-XeL groups	**t-test** t: 0.1577; degrees of freedom: 23.955; p-value: 0.876; 95% Confidence Interval: [-0.8636696 1.0065268] No significant difference found
Responses to the statement ”If I need to learn about new Grid technologies in future I would use the eLGrid eLearning system” for XeL and non-XeL groups	**t-test** t: -1.0647; degrees of freedom: 21.971; p-value: 0.2986; 95% Confidence Interval: [-1.4389466 0.4627561] No significant difference found
Responses to the statement ”The theory in this course complemented and supported the practical exercises” for XeL and non-XeL groups	**t-test** t: -2.3671; degrees of freedom: 24.808; p-value: 0.02604; 95% Confidence Interval: [0.1089553 1.5723634] Significant difference found at alpha = 0.05

Continued on next page

Table L.3 – continued from previous page

Variable	**Results**
Responses to the statement ”I understood how my test answers affected the traffic-light indicators” for XeL and non-XeL groups	**t-test** t: 1.5749 degrees of freedom: 23.744 p-value: 0.1285 95% Confidence Interval: [-0.1852771 1.3757533] No significant difference found
Responses by the XeL group to the statements ”I understood how the results of my practical exercises affected the traffic-light indicators” and ”I understood how my test answers affected the traffic-light indicators”	**t-test** t: 0.7521 degrees of freedom: 21.942 p-value: 0.46 95% Confidence Interval: [-0.4394301 0.9394301] No significant difference found

Continued on next page

Table L.3 – continued from previous page

Variable	**Results**
Responses to the statement ”I understood how the results of my practical exercises affected the traffic-light indicators” by the XeL group and to the statement ”I understood how my test answers affected the traffic-light indicators” by the non-XeL group	**t-test** t / degrees of freedom / p-value / 95% Confidence Interval 2.2822 / 23.45 / 0.03186 / [0.07988857 1.61058762] Significant difference found at alpha = 0.05
Responses to the statement ”I understood how my test answers affected the traffic-light indicators” for XeL and non-XeL groups	**t-test** t / degrees of freedom / p-value / 95% Confidence Interval 1.5749 / 23.744 / 0.1285 / [-0.1852771 1.3757533] No significant difference found
Responses to the statement ”The tests were a good way for eLGrid to evaluate my knowledge” for XeL and non-XeL groups	**t-test** t / degrees of freedom / p-value / 95% Confidence Interval -0.0345 / 22.101 / 0.9728 / [-0.7279446 0.7041351] No significant difference found

Continued on next page

Table L.3 – continued from previous page

Variable	**Results**
Responses to the statement ”The practical exercises were a good way for me to evaluate my knowledge” for XeL and non-XeL groups	**t-test** t: -0.2109; degrees of freedom: 23.991; p-value: 0.8347; 95% Confidence Interval: [-0.6419553 0.5229077] No significant difference found
Responses by the XeL group to the statements ”The practical exercises were a good way for eLGrid to evaluate my knowledge” and ”The tests were a good way for eLGrid to evaluate my knowledge”	**t-test** t: 0.6106; degrees of freedom: df = 22; p-value: 0.5477; 95% Confidence Interval: [-0.3993713 0.7327047] No significant difference found
Responses to the statements ”The practical exercises were a good way for eLGrid to evaluate my knowledge” by the XeL group and ”The tests were a good way for eLGrid to evaluate my knowledge” by the non-XeL group	**t-test** t: 0.4481; degrees of freedom: 22.101; p-value: 0.6584; 95% Confidence Interval: [-0.5612779 0.8708018] No significant difference found

Continued on next page

Table L.3 – continued from previous page

Variable	**Results**			
	t-test			
	t	degrees of freedom	p-value	95% Confidence Interval
Responses to the statements ”The fact that the traffic-light indicators took into account the results of the practical exercises as well as my test answers was helpful to me” by the XeL group and ”It would be useful if the traffic-light indicators took into account the results of my practical exercises as well as my test answers” by the non-XeL group	0.513	19.443	0.6137	[-0.6585811 1.0871525]
	No significant difference found			

Appendix M

Sample post-course survey

Learner Survey

In order to help improve the eLGrid eLearning system we would appreciate it if you could take a few moments to fill in the following questionnaire. For each statement please indicate with a tick (see example below) how much you agree or disagree on a scale of 1 to 5 where 1 means you strongly disagree and 5 means you strongly agree.

Name: ____

WebCom

1. The course expected too much prior knowledge

Strongly Disagree	1	2	3	4	5	Strongly Agree

2. The course was too simple, explaining things that I already knew

Strongly Disagree	1	2	3	4	5	Strongly Agree

3. The practical instructions were easy to understand

Strongly Disagree	1	2	3	4	5	Strongly Agree

4. The practical environment (WebCom IDE) was easy to access

Strongly Disagree	1	2	3	4	5	Strongly Agree

5. The practical exercises helped me to understand and apply the course concepts

Strongly Disagree	1	2	3	4	5	Strongly Agree

6. The theory in this course supported and complemented the practical exercises

Strongly Disagree	1	2	3	4	5	Strongly Agree

7. The presence of a tutor was necessary for me to complete this course, I would have had trouble if I was trying to use this course on my own

Strongly Disagree	1	2	3	4	5	Strongly Agree

Overall Impressions

8. The eLGrid system is easy to use

Strongly Disagree	1	2	3	4	5	Strongly Agree

9. The practical environment (WebCom IDE) is easy to use

Strongly Disagree	1	2	3	4	5	Strongly Agree

10. The traffic-light indicators helped me to navigate the courses

Strongly Disagree	1	2	3	4	5	Strongly Agree

11. I understood how my test answers affected the traffic-light indicators

Strongly Disagree	1	2	3	4	5	Strongly Agree

12. I understood how the results of the final practical exercise affected the traffic-light indicators

Strongly Disagree	1	2	3	4	5	Strongly Agree

13. The tests were a good way for the system to evaluate my knowledge

Strongly Disagree	1	2	3	4	5	Strongly Agree

14. The practical exercises were a good way for me to evaluate my knowledge

Strongly Disagree	1	2	3	4	5	Strongly Agree

15. The final practical exercise was a good way for the system to evaluate my knowledge

Strongly Disagree	1	2	3	4	5	Strongly Agree

16. The fact that the traffic-light indicators took into account the results of the final practical exercise as well as my test answers was helpful to me

Strongly Disagree	1	2	3	4	5	Strongly Agree

17. The instant feedback on the results of the final practical exercise was helpful

Strongly Disagree	1	2	3	4	5	Strongly Agree

18. In the future I will probably use the technologies taught in this course (WebCom)

Strongly Disagree	1	2	3	4	5	Strongly Agree

19. If I need to learn about Grid technologies in future I would use the eLGrid eLearning system

Strongly Disagree	1	2	3	4	5	Strongly Agree

20. Please include any other comments on the eLGrid eLearning system below. Thank you for taking the time to complete the user survey!

References

[1] Enabling Grids for E-SciencE (EGEE, EGEE-II, EGEE-III). EU Framework Programme research projects numbers INFSO-RI-508833, INFSO-RI-222667, INFSO-RI-031688, 2004 – 2010. URL `http://www.eu-egee.org/`, last access 20/04/2011.

[2] International Collaboration to Extend and Advance Grid Education (ICEAGE). EU Framework Programme research project number INFSO-SSA-26637, 2006 – 2008. URL `http://www.iceage-eu.org/`, last access 24/04/2011.

[3] e-Infrastructures Reflection Group (e-IRG), 2011. URL `http://www.e-irg.eu/`, last access 21/04/2011.

[4] OGF Education and Training Community Group (ET-CG). Open Grid Forum (OGF), 2008. URL `https://forge.gridforum.org/sf/projects/et-cg`, last access 11/04/2011.

[5] e-IRG White Paper. Technical report, e-Infrastructures Reflection Group (e-IRG), 2009. URL `http://www.e-irg.eu/images/stories/e-irg_white_paper_ii_web1.pdf`, last access 24/04/2011.

[6] Malcolm Atkinson and Elizabeth Vander Meer. Policy for Supporting Grid and e-Science Education and Training. Technical report, Open Grid Forum (OGF), 2009. URL `http://www.gridforum.org/documents/GFD.153.pdf`, last access 25/04/2011.

[7] Education and Training Task Force Report. Technical report, e-Infrastructures Reflection Group (e-IRG), 2008. URL http://www.e-irg.eu/images/stories/publ/task_force_reports/ettf_long_report_final_july08.pdf, last access 24/04/2011.

[8] Gráinne Conole. Mediating artefacts to guide choice in creating and undertaking learning activities. Discussion paper for the Open University Computers and Learning Research Group (CALRG) Seminar, Open University, UK, November 2005.

[9] Peter Brusilovsky. Adaptive Hypermedia. *User Modeling and User-Adapted Interaction*, 11(1):87 – 110, 2001. URL http://dx.doi.org/10.1023/A:1011143116306, last access 25/04/2011.

[10] Kathryn Cassidy, Jason McCandless, Stephen Childs, John Walsh, Brian Coghlan, and Declan Dagger. Combining a virtual grid testbed and grid eLearning courseware. In *Proceedings Cracow Grid Workshop 2006 (CGW06)*, pages 50 – 57, Cracow, Poland, October 2006. Academic Computer Centre CYFRONET AGH.

[11] Brian Coghlan, John Walsh, and David O.Callaghan. Grid-Ireland Deployment Architecture. In *Proceedings 3rd European Grid Conference (EGC'2005)*, pages 354 – 363, Amsterdam, February 2005. Springer-Verlag. URL http://www.scss.tcd.ie/David.OCallaghan/grid/grid-ireland/architecture.pdf, last accessed 28/04/2011.

[12] Introduction to Developing Grid Applications. Workshop, Trinity College Dublin, 2006. URL http://grid.ie/grid-event-2006/march-course/index.php, last access 25/04/2011.

[13] Linda A. Cornwall, Jens Jensen, David P. Kelsey, Ákos Frohner, Daniel Kouřil, Franck Bonnassieux, Sophie Nicoud, Károly Lőrentey, Joni Hahkala, Mika Silander, Roberto Cecchini, Vincenzo Ciaschini, Luca dell'Agnello, Fabio Spataro, David O'Callaghan, Olle Mulmo, Gian Luca Volpato, David

Groep, Martijn Steenbakkers, and Andrew McNab. Authentication and Authorization Mechanisms for Multi-Domain Grid Environments. *Journal of Grid Computing*, 2(4):301 – 311, 2004. URL `http://dx.doi.org/10.1007/s10723-004-8182-y`, last access 24/04/2011.

[14] Peter Brusilovsky. Adaptive Navigation Support: From Adaptive Hypermedia to the Adaptive Web and Beyond. *PsychNology Journal*, 2(1):7 – 23, 2004.

[15] European Grid Infrastructure (EGI), 2011. URL `http://www.egi.eu/`, last access 25/04/2011.

[16] Hermann Ebbinghaus. *Memory : a contribution to experimental psychology / Translated by Henry A. Ruger and Clara E. Bussenius, and with a new introd. by Ernest R. Hilgard*. Dover Publications, New York, 1964.

[17] Nicholas Cepeda, Noriko Coburn, Doug Rohrer, J T Wixted, Mike Mozer, and Harold Pashler. Optimizing distributed practice: Theoretical analysis and practical implications. *Experimental Psychology*, 56(4):236 – 246, 2009. URL `http://escholarship.org/uc/item/4hc333j8`, last access 25/04/2011.

[18] Ruth Colvin Clark and Richard E. Mayer. *E-Learning and the science of instruction : proven guidelines for consumers and designers of multimedia learning*. Jossey-Bass/Pfeiffer, San Francisco, CA, US, first edition, 2003.

[19] Diego R. Lopez, José-Manuel Macías, Maurizio Molina, Jürgen Rauschenbach, Andreas Åkre Solberg, Manuela Stanica, and GN2 JRA5 team. Deliverable DJ5.2.3,3: Best Practice Guide — AAI Cookbook. Technical report, GÉANT, 2008. URL `http://www.geant2.net/upload/pdf/GN2-08-130-DJ5-2-3-3_eduGAIN_AAI_CookBook-1.pdf`, last access 28/04/2011.

[20] Lauren B. Resnick and Daniel P. Resnick. Assessing the Thinking Curriculum: New Tools for Educational Reform. In Bernard R. Gifford and Mary Catherine O'Connor, editors, *Changing Assessments: Alternative Views of Aptitude, Achievement and Instruction*, pages 37 – 76. Kluwer Academic Publishers, Boston, MA, 1992.

[21] Dale H. Schunk. *Learning Theories: An Educational Perspective*. Prentice Hall, Harlow, fifth edition, 2008.

[22] Burrhus Frederic Skinner. *About Behaviorism*. Random House, Vintage, New York, 1976.

[23] Abram Amsel. *Behaviorism, neobehaviorism, and cognitivism in learning theory: historical and contemporary perspectives*. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1988.

[24] Svein Sjøberg. Constructivism and learning. In Penelope Peterson, Eva Baker, and Barry McGaw, editors, *International Encyclopedia of Education*, pages 485 – 490, Oxford, May 2010. Elsevier Science.

[25] Charles C. Bonwell and James A. Eison. Active Learning: Creating Excitement in the Classroom. Technical report, ERIC Higher Education Digests, Washington DC., September 1991. URL `http://www.ntlf.com/html/lib/bib/91-9dig.htm`, last access 27/04/2011.

[26] Richard E. Mayer. Should there be a three-strikes rule against pure discovery learning? The case for guided methods of instruction. *American Psychologist*, 59(1):14–19, 2004.

[27] Howard S. Barrows and Robyn M. Tamblyn. *Problem-based learning: an approach to medical education*. Springer, New York, 1980.

[28] John Canavan. Personalised E-Learning Through Learning Style Aware Adaptive Systems. Master's thesis, Trinity College Dublin, 2004. URL `http://www.scss.tcd.ie/publications/tech-reports/reports.05/TCD-CS-2005-08.pdf`, last access 27/04/2011.

[29] Rohani A. Tarmizi and John Sweller. Guidance during mathematical problem solving. *Journal of Educational Psychology*, 80(4):424 – 436, 1988. URL `http://search.ebscohost.com/login.aspx?direct=true&db=pdh&AN=edu-80-4-424&site=ehost-live`, last access 24/04/2011.

[30] Tim O'Reilly. Hard Work and Practice in Programming. *O'Reilly Radar*, December 2008. e-Zine URL "http://radar.oreilly.com/2008/12/hard-work-and-practice-in-programming.html", last access 24/04/2011.

[31] Nicholas Cepeda, Harold Pashler, Edward Vul, John T. Wixted, and Doug Rohrer. Distributed practice in verbal recall tasks: A review and quantitative synthesis. *Psychological Bulletin*, 132(3):354 – 380, 2006.

[32] John Sweller. Implications of Cognitive Load Theory for Multimedia Learning. In Richard Mayer, editor, *The Cambridge handbook of multimedia learning*, pages 19 – 30. Cambridge University Press, New York, first edition, 2005.

[33] Roxana Moreno and Richard E. Mayer. Cognitive principles of multimedia learning: The role of modality and contiguity. *Journal of Educational Psychology*, 91:358 – 368, 1999.

[34] Slava Kalyuga. Expertise Reversal Effect and Its Implications for Learner-Tailored Instruction. *Educational Psychology Review*, 19(4):509 – 539, 2007. URL http://dx.doi.org/10.1007/s10648-007-9054-3, last access 24/04/2011.

[35] Paul Nicholson. A History of E-Learning. In Baltasar Fernández-Manjón, Juan Manuel Sánchez-Pérez, Juan Antonio Gómez-Pulido, Miguel Angel Vega-Rodríguez, and José Bravo-Rodríguez, editors, *Computers and Education*, pages 1 – 11. Springer Netherlands, 2007. URL http://dx.doi.org/10.1007/978-1-4020-4914-9_1, last access 24/04/2011.

[36] Blackboard LMS. URL http://www.blackboard.com/, last access 25/04/2011.

[37] Moodle LMS. URL http://moodle.org/, last access 25/04/2011.

[38] Sakai LMS. URL http://sakaiproject.org/, last access 25/04/2011.

[39] Robert Cosgrave, Angelica Risquez, Theresa Logan-Phelan, Thomas Farrelly, Eamon Costello, Marion Palmer, Claire McAvinia, Nuala Harding, and Noreen

Vaughan. Usage and uptake of Virtual Learning Environments in Ireland: Findings from a multi-institutional study. *All Ireland Journal of Teaching and Learning in Higher Education (AISHE-J)*, 3(1):30.1 – 30.14, 2011.

[40] Pierre Dillenbourg, Daniel Schneider, and Paraskevi Synteta. Virtual Learning Environments. In *Proceedings of the 3rd Congress on Information and Communication Technologies in education*, pages 3 – 18, Rhodes, Greece, 2002. Kastaniotis Editions. URL http://hal.archives-ouvertes.fr/hal-00190701/en/, last access 25/04/2011.

[41] Rubicon LMS. URL http://www.intuition.com/, last access 25/04/2011.

[42] SumTotal LMS. URL http://www.sumtotalsystems.com/, last access 25/04/2011.

[43] Marcus Specht, Dominique Verpoorten, Christian Glahn, Milos Kravcik, and Stefaan Ternier. *Personalisation of Learning in Virtual Learning Environments*, volume 5794/2009 of *Lecture Notes in Computer Science*, pages 52 – 66. Springer Berlin / Heidelberg, 2009. URL http://dx.doi.org/10.1007/978-3-642-04636-0, last access 25/04/2011.

[44] Dongming Xu and Huaiqing Wang. Intelligent agent supported personalization for virtual learning environments. *Decision Support Systems*, 42(2):825 – 843, 2006. URL http://linkinghub.elsevier.com/retrieve/pii/S0167923605000989, last access 25/04/2011.

[45] Micheál Tiarnaigh. Adaptive Moodle. Final Year Project, Trinity College Dublin, 2005. URL http://www.scss.tcd.ie/undergraduate/bacsll/bacsll_web/tiarnaigh0405.pdf, last access 25/04/2011.

[46] Ergin Murat Altuner and Mustafa Ali Türker. Learning Designs Supporting Localisation for Personalised and Adaptive E-Learning. In *Proceedings Workshops of the Fourth International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH2006)*, Dublin, Ireland, June 2006. URL http://hdl.handle.net/1820/749, last access 28/04/2011.

[47] Chih-Ming Chen, Hahn-Ming Lee, and Ya-Hui Chen. Personalized e-Learning system using Item Response Theory. *Computers & Education*, 44(3):237 – 255, April 2005. URL http://portal.acm.org/citation.cfm?id=1066362.1066365, last access 25/04/2011.

[48] Neil D. Fleming. *Teaching and Learning Styles: VARK Strategies*. Privately Published, Neil D. Flemming, Christchurch, New Zealand, fifth edition, 2001.

[49] Cristina Hava Muntean and Jennifer McManis. End-User Quality of Experience Oriented Adaptive E-learning System. *Journal of Digital Information*, 7(1), 2006. URL http://journals.tdl.org/jodi/article/view/179, last access 28/04/2011.

[50] Paul De Bra, Peter Brusilovsky, and Geert-Jan Houben. Adaptive hypermedia: from systems to framework. *ACM Computing Surveys*, 31(4), December 1999. URL http://doi.acm.org/10.1145/345966.345996, last access 25/04/2011.

[51] Peter Honey and Alan Mumford. *The Manual of Learning Styles*. Privately Published, Peter Honey Publications, Berkshire, third edition, 1992.

[52] Walter L. Leite, Marilla Svinicki, and Yuying Shi. Attempted validation of the scores of the VARK: Learning styles inventory with multitrait–multimethod confirmatory factor analysis models. *Educational and Psychological Measurement*, 70(2):323 – 339, 2010. URL http://epm.sagepub.com/content/70/2/323.abstract, last access 27/04/2011.

[53] David A. Kolb. *Experiential learning — Experience as the Source of Learning and Development*. Prentice-Hall, New Jersey, 1984.

[54] Dessislava Vassileva. Storyboard Design for Adaptive E-learning Based on Learning Styles. In *Proc. Second International Conference on Software, Services and Semantic Technologies (S3T)*, Varna, Bulgaria, September 2010. URL http://hdl.handle.net/10506/389, last access 28/04/2011.

[55] Howard Gardner and Thomas Hatch. Educational Implications of the Theory of Multiple Intelligences. *Educational Researcher*, 18(8):4 – 10, 1989. URL `http://edr.sagepub.com/content/18/8/4.abstract`, last access 25/04/2011.

[56] Gordon Pask. Styles and strategies of learning. *British Journal of Educational Psychology*, 46(2):128 – 148, 1976.

[57] Liam Hudson. *Contrary imaginations*. Methuen, London, 1966.

[58] Elvira Popescu, Philippe Trigano, and Costin Badica. Towards a Unified Learning Style Model in Adaptive Educational Systems. In *Proc. Seventh IEEE International Conference on Advanced Learning Technologies (ICALT 2007)*, pages 804 – 808, Niigata, Japan, July 2007. IEEE Computer Society.

[59] K. A. Papanikolaou and M. Grigoriadou. Accommodating learning style characteristics in Adaptive Educational Hypermedia. In *Proceedings Workshop on Individual Differences in Adaptive Hypermedia in Third International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH2004), Part I*, pages 77 – 86, Eindhoven, Netherlands, 2004.

[60] Andreea Molnar and Cristina Hava Muntean. Performance Aware and Cost Oriented Adaptive e-Learning Framework. In *Proceedings IADIS International Conference e-Learning*, Portugal, 2009. International Association for Development of the Information Society (IADIS). URL `http://pel.ucd.ie/node/261`, last access 28/04/2011.

[61] Anastasios A. Economides. Adaptive context-aware pervasive and ubiquitous learning. *International Journal of Technology Enhanced Learning*, 1(3):169 – 192, 2009.

[62] Intuition Publishing Ltd. URL `http://www.intuition.com/`, last access 25/04/2011.

[63] Emmanuel Blanchard, Ryad Razaki, and Claude Frasson. Cross-Cultural Adaptation of eLearning Contents: a Methodology. In *Proc. International Confer-*

ence on E-Learning (ELEARN2005), pages 1895 – 1902, Vancouver, Canada, 2005.

[64] Owen Conlan and Vincent Wade. Evaluation of APeLS — An Adaptive eLearning Service based on the Multi-model, Metadata-driven Approach. In *Proceedings Third International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, Eindhoven, The Netherlands, 2004.

[65] Clark Aldrich. *The Complete Guide to Simulations and Serious Games*. Pfeiffer, San Francisco, CA, 2009.

[66] Mary Ulicsak and Martha Wright. Games in education: Serious Games. Technical report, Futurelab, Bristol, 2010. URL http://media.futurelab.org.uk/resources/documents/lit_reviews/Serious-Games_Review.pdf, last access 27/04/2011.

[67] Adobe Captivate. URL http://www.adobe.com/products/captivate/, last access 25/04/2011.

[68] Assima. URL http://www.assima.net/, last access 25/04/2011.

[69] RapidBuilder. URL http://www.xstreamsoftware.com/rapidbuilder_description.htm, last access 25/04/2011.

[70] Raptivity. URL http://www.raptivity.com/, last access 25/04/2011.

[71] SkillSoft. URL http://www.skillsoft.com/products/course_collections/skillsims.asp, last access 25/04/2011.

[72] Jeremy Jones. Vivio — A System for Creating Interactive Reversible E-Learning Animations for the WWW. In *Proceedings IEEE Symposium on Visual Languages and Human-Centric Computing*, pages 131 – 133, Rome, Italy, 2004. IEEE Computer Society.

[73] Conor Gaffney, Declan Dagger, and Vincent Wade. Evaluation of ACTSim: A Composition Tool for Authoring Adaptive Soft Skill Simulations. In Wolfgang Nejdl, Judy Kay, Pearl Pu, and Eelco Herder, editors, *Proceedings Fifth*

International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH2008), volume 5149 of *Lecture Notes in Computer Science*, pages 113 – 122. Springer Berlin / Heidelberg, 2008. URL http://dx.doi.org/10.1007/978-3-540-70987-9_14, last access 25/04/2011.

[74] Joanna Bull. Using Technology to Assess Student Learning. Technical report, Project ALTER, 1993. URL http://www.caacentre.ac.uk/dldocs/alter.pdf, last access 25/04/2011.

[75] C Leacock and M Chodorov. Scoring Free-Responses Automatically: A Case Study of a Large-Scale Assessment. *Computers and Humanities*, 37(4):389 – 405, 2003. URL http://www.ets.org/Media/Research/pdf/erater_examens_leacock.pdf, last access 27/04/2011.

[76] Case study 5: Making the most of a computer-assisted assessment system. Technical report, JISC, UK, 2010. URL http://www.jisc.ac.uk/media/documents/programmes/elearning/digiassess_makingthemost.pdf, last access 20/04/2011.

[77] James McCarthy, John Wayne, and John Morris. Closed-Loop Adaptive Education. In *Proceedings of World Conference on Educational Multimedia, Hypermedia and Telecommunications*, pages 1268 – 1273, Chesapeake, VA, USA, 2001.

[78] David Fergusson, Richard Hopkins, Diego Romano, Elizabeth Vander Meer, and Malcolm Atkinson. Distributed Computing Education, Part 2: International Summer Schools. *IEEE Distributed Systems Online*, 9(7), July 2008. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4589948, last access 28/04/2011.

[79] Petar Jandric, Morgane Artacho, Richard Hopkins, and David Fergusson. From distributed computing to distributed learning. In *Proceedings of the Seventh IASTED International Conference on Web-based Education (WBE '08)*, pages

127 – 132, Innsbruck, Austria, 2008. ACTA Press. URL http://portal.acm.org/citation.cfm?id=1722815.1722843, last access 25/04/2011.

[80] GridKa School. URL http://gridka-school.scc.kit.edu/, last access 25/04/2011.

[81] Ruth Pordes, Mine Altunay, Paul Avery, Alina Bejan, Kent Blackburn, Alan Blatecky, Rob Gardner, Bill Kramer, Miron Livny, John McGee, Maxim Potekhin, Rob Quick, Doug Olson, Alain Roy, Chander Sehgal, Torre Wenaus, Mike Wilde, and Frank Würthwein. New science on the Open Science Grid. *Journal of Physics: Conference Series*, 125(1), 2008. URL http://stacks.iop.org/1742-6596/125/i=1/a=012070, last access 25/04/2011.

[82] EU Framework 7 Funding Programme. URL http://cordis.europa.eu/fp7/home_en.html, last access 25/04/2011.

[83] DCI Summer School. Budapest, Hungary, July 2011. URL http://www.lpds.sztaki.hu/eudciss2011/, last access 25/04/2011.

[84] Biomed Grid School, May 2008. URL http://www.itb.cnr.it/indico/conferenceDisplay.py?confId=5, last access 25/04/2011.

[85] CERN School of Computing. URL https://csc.web.cern.ch/CSC/, last access 25/04/2011.

[86] MSc in Grid Computing. University of Amsterdam. URL http://studiegids.uva.nl/web/uva/sgs/en/p/178.html, last access 25/04/2011.

[87] MSc in Distributed Scientific Computing. University of Edinburgh. URL http://www.nesc.ac.uk/msc/, last access 25/04/2011.

[88] Maozhen Li and Marios Hadjinicolaou. Introducing Grid computing to a masters course: the design rationale and practice. In *Proceedings of the 5th WSEAS/IASME international conference on Engineering education*, pages 150 – 154, Heraklion, Greece, 2008. World Scientific and Engineering Academy

and Society (WSEAS). URL http://portal.acm.org/citation.cfm?id=1581120.1581151, last access 25/04/2011.

[89] e-IRG Education and Training Wiki. URL https://eirgsp-wiki.grnet.gr/bin/view/Main/TrainingAndEducation, last access 25/04/2011.

[90] CS4021 Advanced Computer Architecture course at Trinity College Dublin, 2011. URL https://www.scss.tcd.ie/undergraduate/ba/assets/pdf/modules/CS4021_0910.pdf, last access 25/04/2011.

[91] Open Grid Forum (OGF), 2006 – 2011. URL http://www.ogf.org/, last access 21/04/2011.

[92] Rüdiger Berlich, Kilian Schwarz, and Kathryn Cassidy. Towards Professional Grid Certification. Technical report, OGF, 2009. URL http://www.gridforum.org/documents/GFD.155.pdf, last access 25/04/2011.

[93] Curricula for Undergraduate and Masters Level Courses in e-Science: Report from the ICEAGE Curricula Development Workshop. Technical report, EU ICEAGE project, Brussels, 2008. URL http://www.iceage-eu.org/v2/documents/Curricula%20for%20e-Science%20Education.doc, last access 28/04/2011.

[94] Boon Low, Kathryn Cassidy, David Fergusson, Malcolm Atkinson, Elizabeth Vander Meer, and Mags McGeever. Distributed Computing Education, Part 5: Coming to Terms with Intellectual Property Rights. *IEEE Distributed Systems Online*, 9(12), December 2008. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4755177, last access 28/04/2011.

[95] Grid People Registry. EU ICEAGE project. URL http://www.iceage-eu.org/learnerRegistry/, last access 25/04/2011.

[96] e-IRG Roadmap 2010. Technical report, e-Infrastructures Reflection Group (e-IRG), 2010. http://www.e-irg.eu/images/stories/eirg_roadmap_2010_layout_final.pdf.

[97] European Middleware Initiative (EMI). EU Framework Programme research project number INFSO-RI-261611, 2011. URL `http://www.eu-emi.eu/en/`, last access 25/04/2011.

[98] Paula Eerola, Tord Ekelöf, Mattias Ellert, Michael Grønager, John Renner Hansen, Sigve Haug, Josva Kleist, Aleksandr Konstantinov, Balázs Kónya, Farid Ould-Saada, Oxana Smirnova, Ferenc Szalai, and Anders Wäänänen. Roadmap for the ARC grid middleware. In *Proceedings of the 8th international conference on Applied parallel computing: state of the art in scientific computing*, PARA'06, pages 471 – 479, Umeå, Sweden, 2007. Springer-Verlag. URL `http://portal.acm.org/citation.cfm?id=1775059.1775126`, last access 25/04/2011.

[99] Patrick Fuhrmann and Volker Gülzow. dCache, Storage System for the Future. In Wolfgang Nagel, Wolfgang Walter, and Wolfgang Lehner, editors, *Euro-Par 2006 Parallel Processing*, volume 4128 of *Lecture Notes in Computer Science*, pages 1106–1113. Springer Berlin / Heidelberg, 2006. URL `http://dx.doi.org/10.1007/11823285_116`, last access 25/04/2011.

[100] Massimo Lamanna, Birger Koblitz, T Chen, W Ueng, J Herrala, D Liko, A Maier, J Moscicki, A Peters, F Orellana, V Pose, A Demichev, and Derek Feichtinger. Experiences with the gLite Grid Middleware. In *Proceedings International Conference on Computing in High Energy and Nuclear Physics (CHEP)*, Interlaken, Switzerland, 2004. URL `http://indico.cern.ch/contributionDisplay.py?contribId=305&sessionId=7&confId=0`, last access 28/04/2011.

[101] Dietmar W. Erwin. UNICORE — a Grid computing environment. *Concurrency and Computation: Practice and Experience*, 14(13 – 15):1395 – 1410, 2002. URL `http://dx.doi.org/10.1002/cpe.691`, last access 25/04/2011.

[102] Adobe Connect. URL `http://www.adobe.com/products/adobeconnect.html`, last access 25/04/2011.

[103] StratusLab. EU Framework Programme research project number INFSO-RI-261552, 2011. URL http://www.stratuslab.eu/, last access 21/04/2011.

[104] StratusLab User Tutorial Video. StratusLab project, 2010. URL http://www.youtube.com/watch?v=eR8OgKsuBoo, last access 25/04/2011.

[105] e-Science Talk. URL http://www.e-sciencetalk.org/, last access 25/04/2011.

[106] International Science Grid This Week. URL http://www.isgtw.org/, last access 25/04/2011.

[107] GridGuide. URL http://www.gridguide.org/, last access 25/04/2011.

[108] GridCafe. URL http://www.gridcafe.org/, last access 25/04/2011.

[109] GridCast. URL http://gridcast.web.cern.ch/Gridcast/, last access 25/04/2011.

[110] GridBriefings. URL http://www.e-sciencetalk.org/briefings.php, last access 25/04/2011.

[111] Putting the 'e' in education: eLearning and grid computing. e-Science Talk GridBriefing, July 2010. URL http://www.e-sciencetalk.org/briefings/gridbriefing-eLearning-and-grid-computing.pdf, last access 25/04/2011.

[112] EMI 2nd in-reach training event (online). European Middleware Initiative (EMI), February – March 2011. URL http://indico.cern.ch/conferenceDisplay.py?confId=127278, last access 25/04/2011.

[113] The GILDA Training Wiki. URL https://grid.ct.infn.it/twiki/bin/view/GILDA/WebHome, last access 25/04/2011.

[114] OSG Education, Outreach and Training Wiki. Open Science Grid. URL https://twiki.grid.iu.edu/bin/view/Education/WebHome, last access 25/04/2011.

[115] Clinton S. Potter, Bridget Carragher, Liana Carroll, Charles Conway, Benjamin Grosser, Janet Hanlon, Nick Kisseberth, Scott Robinson, Umesh Thakkar, and Daniel Weber. Bugscope: A Practical Approach to Providing Remote Microscopy for Science Education Outreach. *Microscopy and Microanalysis*, 7:249 – 252, 2001.

[116] Ioannis T. Christou, Thanassis Tiropanis, Sofia Tsekeridou, and Konstantinos Roussos. Grid-Based Interactive Virtual Scientific Experiments for Distributed Virtual Communities. *International Journal of Continuing Engineering Education and Life-Long Learning*, 20(6):443 – 464, 2010. URL http://eprints.ecs.soton.ac.uk/15273/, last access 25/04/2011.

[117] Creative Commons. URL http://creativecommons.org/, last access 25/04/2011.

[118] Joachim Götze, Bernd Reuther, Paul Müller, Viktor Achter, Sebastian Breuers, Marc Seifert, and Ulrich Lang. SuGI — Portal and Training Systems for Grid Middlewares. In *Proceedings of the 2009 35th Euromicro Conference on Software Engineering and Advanced Applications (SEAA '09)*, pages 585 – 591, Washington, DC, USA, 2009. IEEE Computer Society. URL http://dx.doi.org/10.1109/SEAA.2009.54, last access 25/04/2011.

[119] Lisa Childers, Terry Disz, Robert Olson, Michael E. Papka, Rick Stevens, and Tushar Udeshi. Access Grid: Immersive Group-to-Group Collaborative Visualization. Argonne National Laboratory, 2000. URL http://ftp.mcs.anl.gov/pub/tech_reports/reports/P821.doc, last accessed 28/04/2011.

[120] Hsin-Chuan Ho, Chao-Tung Yang, and Chi-Chung Chang. Building an e-Learning platform by Access Grid and Data Grid technologies. In *Proc. e-Technology, e-Commerce and e-Service, 2004 IEEE International Conference on*, pages 452 – 455, Taipei, Taiwan, March 2004.

[121] Hyunju Kim, Loretta A. Moore, and Geoffrey Fox. An Experience on a Distance Education Course over the Access Grid Nodes. In *Proc. Fourth*

International Conference on Education and Information Systems, Technologies and Applications (EISTA), Orlando, Florida, July 2006. URL http://grids.ucs.indiana.edu/ptliupages/publications/eista06.pdf, last access 28/04/2011.

[122] Michalis Xenos, Bill Vassiliadis, and Athanassios Skodras. GRID Technologies => 'Education' = 'Distance Education'. In *Proceedings 1st International ELeGI Conference on Advanced Technology for Enhanced Learning*, Vico Equense (Naples) Italy, 2005. URL http://telearn.archives-ouvertes.fr/hal-00190564/PDF/Xenos-M_GRID_2005.pdf, last access 25/04/2011.

[123] Geoffrey Fox. Architecture and Implementation of a Collaborative Computing and Education Portal. Technical report, U.S. Army Corps of Engineers Waterways Experiment Station (ERDC), May 2001. URL http://citeseer.ist.psu.edu/viewdoc/download?doi=10.1.1.23.6159&rep=rep1&type=pdf, last accessed 28/04/2011.

[124] Bill Vassiliadis, Michalis Nik Xenos, and Antonia Stefani. Enabling Enhanced Learning in an Open University Environment: A Grid-Aware Architecture. *International Journal of Information and Communication Technology Education (IJICTE)*, 5(3):59 – 73, 2009.

[125] European Learning Grid Infrastructure (ELeGI). EU Framework Programme research project number IST-002205. URL http://www.ecs.soton.ac.uk/research/projects/261, last access 25/04/2011.

[126] David Fergusson, Roberto Barbera, Emidio Giorgio, Marco Fargetta, Gergely Sipos, Diego Romano, Malcolm P. Atkinson, and Elizabeth Vander Meer. Distributed Computing Education, Part 4: Training Infrastructure. *IEEE Distributed Systems Online*, 9(10), October 2008. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4752926, last access 28/04/2011.

[127] Roberto Barbera, Leandro Neumann Ciuffo, Emidio Giorgio, Antonio Calanducci, and Valeria Ardizzone. GILDA — Grid INFN Virtual Laboratory for

Dissemination Activities. In *Proc. International Conference on Computer Supported Education (CSEDU)*, pages 273–278, Lisboa, Portugal, 2009.

[128] Instituto Nazionale di Fisica Nucleare (INFN). URL http://www.infn.it/indexen.php, last access 25/04/2011.

[129] GILDA. URL https://gilda.ct.infn.it/, last access 25/04/2011.

[130] R. Barbera, A. Falzone, V. Ardizzone, and D. Scardaci. The GENIUS Grid Portal: Its architecture, improvements of features, and new implementations about authentication and authorization. In *Enabling Technologies: Infrastructure for Collaborative Enterprises, 2007. WETICE 2007. 16th IEEE International Workshops on*, pages 279 – 283, Paris, France, june 2007. IEEE Computer Society.

[131] Csaba Németh, Gábor Dózsa, Róbert Lovas, and Péter Kacsuk. The P-GRADE Grid Portal. In Antonio Laganà, Marina L. Gavrilova, Vipin Kumar, Youngsong Mun, C. J. Kenneth Tan, and Osvaldo Gervasi, editors, *Computational Science and Its Applications – ICCSA 2004*, volume 3044 of *Lecture Notes in Computer Science*, pages 10 – 19. Springer Berlin / Heidelberg, 2004. URL http://dx.doi.org/10.1007/978-3-540-24709-8_2, last access 25/04/2011.

[132] Peter Kacsuk. P-GRADE portal family for grid infrastructures. *Concurrency and Computation: Practice and Experience*, 23(3):235 – 245, 2011. URL http://dx.doi.org/10.1002/cpe.1654, last access 25/04/2011.

[133] EGI: Integrated Sustainable Pan-European Infrastructure for Researchers in Europe (EGI-InSPIRE). EU Framework Programme research project number RI-261323, 2011. URL http://www.egi.eu/projects/egi-inspire/, last access 25/04/2011.

[134] E-Science Grid facility for Europe and Latin America (EELA, EELA-2). EU Framework Programme research projects numbers 026409, 223797, 2006 – 2010. URL http://www.eu-eela.eu/, last access 25/04/2011.

[135] Sustainable e-Infrastructures across Europe and India EU-IndiaGRID, EU-IndiaGrid2. EU Framework Programme research projects numbers RI-031834, RI-246698, 2007 – 2011. URL `http://www.euindiagrid.eu/`, last access 25/04/2011.

[136] EUMEDGRID. EU Framework Programme research project number SSA-026024, 2006 – 2008. URL `http://www2.eumedgrid.eu/`, last access 25/04/2011.

[137] EUChinaGRID. EU Framework Programme research project number SSA-026634, 2006 – 2008. URL `http://www.euchinagrid.org/`, last access 25/04/2011.

[138] Jan Kmuníček and Daniel Kouřil. Central European Grid Infrastructure for Generic Applications. In *Proceedings 2nd Austrian Grid Symposium*, pages 131 – 142, Innsbruck, Austria, 2007. Ossterreichische Computer Gesellschaft.

[139] Greek Research and Technology Network (GRNET). URL `http://www.grnet.gr/`, last access 25/04/2011.

[140] DemoGrid. Globus Alliance, 2011. URL `http://www.globus.org/demogrid/`, last access 11/04/2011.

[141] Vagrant. URL `http://vagrantup.com/`, last access 25/04/2011.

[142] KVM. URL `http://www.linux-kvm.org/`, last access 25/04/2011.

[143] Amazon Elastic Cloud Compute (EC2). URL `http://aws.amazon.com/ec2/`, last access 25/04/2011.

[144] O'Reilly Ellipse, 2011. URL `http://www.oreillyschool.com/courses/ellipse.php`, last access 25/04/2011.

[145] O'Reilly Head First Labs, 2011. URL `http://www.headfirstlabs.com/`, last access 25/04/2011.

[146] Fionnuala O'Donnell. *Simulation frameworks for the teaching and learning of distributed algorithms*. PhD thesis, Trinity College, Dublin, 2006. URL `https://www.scss.tcd.ie/publications/tech-reports/reports.06/TCD-CS-2006-20.pdf`, last access 27/04/2011.

[147] Rajkumar Buyya and Manzur Murshed. GridSim: a toolkit for the modeling and simulation of distributed resource management and scheduling for Grid computing. *Concurrency and Computation: Practice and Experience*, 14(13):1175–1220, 2002.

[148] P Brusilovsky, J Eklund, and E Schwarz. Web-based education for all: a tool for development adaptive courseware. *Computer Networks and ISDN Systems*, 30(1 – 7):291 – 300, 1998. URL `http://linkinghub.elsevier.com/retrieve/pii/S0169755298000828`, last access 01/03/2012.

[149] Paul De Bra, Ad Aerts, Bart Berden, Barend de Lange, Brendan Rousseau, Tomi Santic, David Smits, and Natalia Stash. Aha! the adaptive hypermedia architecture. In *Proceedings of the fourteenth ACM conference on Hypertext and hypermedia*, HYPERTEXT '03, pages 81 – 84, New York, NY, USA, 2003. ACM. URL `http://doi.acm.org/10.1145/900051.900068`, last access 01/03/2012.

[150] Miroslav Ruda, Bob Jones, Patrick Aerts, Willy Weisz, Jacko Koster, David Groep, Michal Turala, and Edward Aronovich. EGI Blueprint. Technical report, EGI, 2008. URL `http://web.eu-egi.eu/documents/other/egi-blueprint/`, last access 25/04/2011.

[151] GANGA. URL `http://ganga.web.cern.ch/ganga/`, last access 25/04/2011.

[152] Migrating Desktop. URL `http://desktop.psnc.pl/`, last access 25/04/2011.

[153] P-Grade. URL `http://portal.p-grade.hu/`, last access 25/04/2011.

[154] EGI-InSPIRE deliverable D6.1, Capabilities Offered by Heavy User Communities. Technical report, EGI-InSPIRE, 2010. URL https://documents.egi.eu/document/154, last access 27/04/2011.

[155] EGI-InSPIRE deliverable D2.3, European Grid Infrastructure – an Integrated Sustainable Pan-European Infrastructure for Researchers in Europe. Technical report, EGI-InSPIRE, 2011. URL https://documents.egi.eu/document/201, last access 25/04/2011.

[156] Yair Levy. Comparing dropouts and persistence in e-Learning courses. *Computers & Education*, 48(2):185 – 204, February 2007. URL http://www.sciencedirect.com/science?_ob=ArticleURL&_udi=B6VCJ-4FDMY9V-3&_user=103681&_rdoc=1&_fmt=&_orig=search&_sort=d&view=c&_acct=C000007920&_version=1&_urlVersion=0&_userid=103681&md5=ce9b3738b849c59c35c647026cf730f8, last access 25/04/2011.

[157] Penjira Kanthawongs. A Structural Model of Student Satisfaction in Web-based ERP-Simulated Learning Environments. *International Journal of Education and Information Technologies*, 5(2):166 – 173, 2011.

[158] William P. Eveland Jr. and Sharon Dunwoody. User Control and Structural Isomorphism or Disorientation and Cognitive Load?: Learning From the Web Versus Print. *Communication Research*, 28(1):48 – 78, 2001. URL http://crx.sagepub.com/cgi/content/abstract/28/1/48, last access 25/04/2011.

[159] Susan Feinberg and Margart Murphy. Applying Cognitive Load Theory to the Design of Web-Based Instruction. In *Proceedings of the 18th annual ACM international conference on Computer documentation: technology & teamwork*, pages 353 – 360, Cambridge, Massachusetts, September 2000. URL http://ieeexplore.ieee.org/iel5/7114/19161/00887293.pdf?arnumber=887293, last access 25/04/2011.

[160] IEEE Standard for Learning Object Metadata. Technical report, IEEE Computer Society, New York, 2002. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1032843, last access 24/04/2011.

[161] IEEE Standard for Learning Technology-Extensible Markup Language (XML) Schema Definition Language Binding for Learning Object Metadata. Technical report, IEEE Computer Society, New York, 2005. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1532505, last access 24/04/2011.

[162] IMS Meta-data Best Practice Guide for IEEE 1484.12.1-2002 Standard for Learning Object Metadata. Technical report, IMS Global Learning Consortium, Inc., August 2006. URL http://www.imsglobal.org/metadata/mdv1p3/imsmd_bestv1p3.html, last access 24/04/2011.

[163] Norm Friesen, Anthony Roberts, and Sue Fisher. Cancore: Metadata for learning objects. *Canadian Journal of Learning and Technology / La revue canadienne de l'apprentissage et de la technologie*, 28(3), 2008. URL http://www.cjlt.ca/index.php/cjlt/article/view/109, last access 28/04/2011.

[164] UK Learning Object Metadata Core. URL http://zope.cetis.ac.uk/profiles/uklomcore, last access 25/04/2011.

[165] Open Grid Services Architecture Glossary of Terms. Technical report, Open Grid Forum (OGF), 2007. URL http://www.ogf.org/documents/GFD.120.pdf, last access 24/04/2011.

[166] OGF Production Grid Infrastructure Glossary of Acronyms and Terms. Technical report, Open Grid Forum (OFG), 2011. URL http://www.gridforum.org/documents/GFD.181.pdf, last access 24/04/2011.

[167] Joan S. Mitchell, Julianne Beall, Giles Martin, Winton E. Matthews, and Gregory R. New, editors. *Dewey decimal classification and relative index / devised by Melvil Dewey*. OCLC Online Computer Library Center, Dublin, Ohio, twenty-second edition, 2003.

[168] Library of Congress Cataloging Distribution Service, editor. *Library of Congress Subject headings*. Library of Congress, Washington(D.C.), seventeenth edition, 1994.

[169] Isabella Peters. *Folksonomies. Indexing and Retrieval in Web 2.0*. De Gruyter Saur, Berlin, first edition, October 2009.

[170] Computer Architecture and Grid Research Group (CAG), 2008. URL https://www.cs.tcd.ie/research_groups/cag/, last access 25/04/2011.

[171] Stephen Childs, Brian Coghlan, and Jason McCandless. GridBuilder: A tool for creating virtual Grid testbeds. In *Proceedings of 2nd IEEE Conference on eScience and Grid computing*, Amsterdam, The Netherlands, December 2006. URL http://hdl.handle.net/2262/16659, last access 28/04/2011.

[172] Brian Coghlan, John Walsh, Geoff Quigley, David O'Callaghan, Stephen Childs, and Eamonn Kenny. Principles of transactional grid deployment. In Peter Sloot, Alfons Hoekstra, Thierry Priol, Alexander Reinefeld, and Marian Bubak, editors, *Advances in Grid Computing — EGC 2005*, volume 3470 of *Lecture Notes in Computer Science*, pages 88 – 97. Springer Berlin / Heidelberg, 2005. URL http://www.springerlink.com/content/3ey8vmgjx1ehgft0/export-citation/, last access 28/04/2011.

[173] Declan Dagger. *Personalised eLearning Development Environments*. PhD thesis, Trinity College Dublin, 2006. URL http://www.scss.tcd.ie/Declan.Dagger/publications/thesis/Declan_Dagger_PhD_Thesis_2006.pdf, last access 27/04/2011.

[174] Dublin Core Metadata Initiative. URL http://dublincore.org/, last access 25/04/2011.

[175] George Lorenzo and John Ittelson. An Overview of e-Portfolios. Technical report, EduCause Learning Initiative, July 2005. URL http://www.sorteoudla.org.mx/promueve/ciedd/CR/tecnologia/AnOverviewofEPortfolios.pdf, last access 25/04/2011.

[176] Roberto Alfieri, Roberto Cecchini, Vincenzo Ciaschini, Luca dell'Agnello, Ákos Frohner, Alberto Gianoli, Károly Lörentey, and Fabio Spataro. VOMS, an Authorization System for Virtual Organizations. In Francisco Fernández Rivera, Marian Bubak, Andrés Gómez Tato, and Ramón Doallo, editors, *Grid Computing*, volume 2970 of *Lecture Notes in Computer Science*, pages 33 – 40. Springer Berlin / Heidelberg, 2004. 10.1007/978-3-540-24689-3_5.

[177] Jim Basney, Marty Humphrey, and Von Welch. The MyProxy online credential repository. *Software: Practice and Experience*, 35(9):801 – 816, 2005.

[178] Synchronized Multimedia Integration Language (SMIL) 1.0 Specification. Technical report, W3C, June 1998. URL `http://www.w3.org/TR/REC-smil/`, last access 25/04/2011.

[179] Andrew W. Cooke, Alasdair J. G. Gray, Werner Nutt, James Magowan, Manfred Oevers, Paul Taylor, Roney Cordenonsi, Rob Byrom, Linda Cornwall, Abdeslem Djaoui, Laurence Field, Steve Fisher, Steve Hicks, Jason Leake, Robin Middleton, Antony J. Wilson, Xiaomei Zhu, Norbert Podhorszki, Brian A. Coghlan, Stuart Kenny, David O'Callaghan, and John Ryan. The Relational Grid Monitoring Architecture: Mediating Information about the Grid. *Journal of Grid Computing*, 2(4):323 – 339, 2004. URL `http://dx.doi.org/10.1007/s10723-005-0151-6`, last access 25/04/2011.

[180] Brian L. Tierney, Ruth Aydt, Dan Gunter, Warren Smith, Martin Swany, Valerie Taylor, and Rich Wolski. A Grid Monitoring Architecture. Technical report, GGF Performance Working Group, 2002. URL `http://www.ogf.org/documents/GFD.7.pdf`, last access 25/04/2011.

[181] Fabrizio Gagliardi, Bob Jones, Mario Reale, and Stephen Burke. European DataGrid Project: Experiences of Deploying a Large Scale Testbed for e-Science Applications. In Maria Calzarossa and Salvatore Tucci, editors, *Performance Evaluation of Complex Systems: Techniques and Tools*, volume 2459 of *Lecture Notes in Computer Science*, pages 255 – 264. Springer Berlin / Heidel-

berg, 2002. URL http://dx.doi.org/10.1007/3-540-45798-4_20, last access 25/04/2011.

[182] Fabrizio Pacini, Alessandro Maraschini, Sara Bertocco, Marco Cecchi, Massimo Sgaravatto, and Alvise Dorigo. Job Description Language Attributes Specification. Technical report, EGEE, 2011. URL https://edms.cern.ch/document/590869/1, last access 25/04/2011.

[183] E Laure, F Hemmer, F Prelz, S Beco, S Fisher, M Livny, L Guy, M Barroso, P Buncic, Peter Z Kunszt, A Di Meglio, A Aimar, A Edlund, D Groep, F Pacini, M Sgaravatto, and O Mulmo. Middleware for the next generation Grid infrastructure. In *Proceedings International Conference on Computing in High Energy and Nuclear Physics (CHEP)*, Interlaken, Switzerland, 2004. URL http://indico.cern.ch/contributionDisplay.py?contribId=247&sessionId=7&confId=0, last access 28/04/2011.

[184] Database Language SQL—Part 2: Foundation (SQL/Foundation). Technical report, ANSI/ISO/IEC International Standard (IS), September 1999. Specification number ISO/IEC 9075-2:1999.

[185] Jack J. Dongarra, Steve W. Otto, Marc Snir, and David Walker. An Introduction to the MPI Standard. Technical report, University of Tennessee, Knoxville, TN, USA, 1995. URL http://citeseer.ist.psu.edu/viewdoc/download?doi=10.1.1.74.3074&rep=rep1&type=pdf, last access 28/04/2011.

[186] John P. Morrison, James J. Kennedy, and David A. Power. WebCom: A Web Based Volunteer Computer. *Journal of Supercomputing*, 18(1):47 – 61, January 2001. URL http://portal.acm.org/citation.cfm?id=370994.371014, last access 25/04/2011.

[187] David A. Power John P. Morrison, Brian Clayton and Adarsh Patil. WebCom-G: Grid Enabled Metacomputing. *The Journal of Neural, Parallel and Scientific Computation. Special Issue on Grid Computing*, 12(3):419 – 438, September

2004. URL http://portal.acm.org/citation.cfm?id=1093534.1093544, last access 25/04/2011.

[188] PIC Simulations of Relativistic Collisionless Shocks. Workshop, Dublin Institute for Advanced Studies (DIAS). URL http://www.dias.ie/picdublin08/index.html, last access 25/04/2011.

[189] EELA-2 Grid School. Workshop, Trinity College Dublin, September 2008. URL http://indico.eu-eela.eu/conferenceDisplay.py?confId=153, last access 25/04/2011.

[190] Kathryn Cassidy, John Walsh, Brian Coghlan, and Declan Dagger. Evaluation of Grid eLearning: eLGrid User Evaluation Experiment Results. In *Proceedings Cracow Grid Workshop 2008 (CGW08)*, pages 365 – 372, Cracow, Poland, October 2008. Academic Computer Centre CYFRONET AGH.

[191] Rensis Likert. A technique for the measurement of attitudes. *Archives of Psychology*, 22(140):1 – 55, 1932.

[192] Susan Jamieson. Likert scales: how to (ab)use them. *Medical Education*, 38(12):1217 – 1218, December 2004. URL http://dx.doi.org/10.1111/j.1365-2929.2004.02012.x, last access 25/04/2011.

[193] Boaz Shulruf, John Hattie, and Robyn Dixon. Factors affecting responses to Likert type questionnaires: introduction of the ImpExp, a new comprehensive model. *Social Psychology of Education*, 11(1):59 – 78, February 2008. URL http://dx.doi.org/10.1007/s11218-007-9035-x, last access 25/04/2011.

[194] Philip M. Podsakoff, Scott B. MacKenzie, Lee Jeong-Yeon, and Nathan P. Podsakoff. Common Method Biases in Behavioral Research: A Critical Review of the Literature and Recommended Remedies. *Journal of Applied Psychology*, 88(5):879, 2003.

[195] Roderick J. A. Little and Donald B. Rubin. The Analysis of Social Science Data with Missing Values. *Sociological Methods & Research*, 18(2 – 3):292 – 326, 1989.

[196] R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2008. URL `http://www.R-project.org`, last access 25/04/2011.

[197] Ian H. Witten and Eibe Frank. *Data Mining: Practical Machine Learning Tools and Techniques*. Elsevier Science, Oxford, second edition, June 2005.

[198] George P. Mccabe and David S. Moore. *Introduction to the Practice of Statistics*. W.H. Freeman & Company, Basingstoke, sixth edition, 2009.

[199] B. L. Welch. The Generalization of 'student's' problem when several different population variances are involved. *Biometrika*, 34(1 – 2):28 – 35, 1947. URL `http://biomet.oxfordjournals.org/content/34/1-2/28.short`, last access 27/04/2011.

[200] Rupert G. Miller. *Simultaneous statistical inference*. Springer-Verlag, New York, second edition, 1985.

[201] Ivy Liu and Alan Agresti. The analysis of ordered categorical data: An overview and a survey of recent developments. *TEST*, 14:1 – 73, 2005.

[202] Bologne Process, 1999. URL `http://ec.europa.eu/education/higher-education/doc1290_en.htm`, last access 25/04/2011.

[203] Leslie Rae. *Assessing the value of your training : the evaluation process from training needs to the report to the board*. Gower, Aldershot, UK, 2002.

[204] Douglas Curran-Everett and Dale J. Benos. Guidelines for reporting statistics in journals published by the American Physiological Society. *Advances in Physiology Education*, 28(3):85 – 87, 2004. URL `http://advan.physiology.org/content/28/3/85.short`, last access 25/04/2011.

[205] Nan L. Maxwell, John R. Mergendoller, and Yolanda Bellisimo. Developing a problem-based learning simulation: An economics unit on trade. *Simulation & Gaming*, 35(4):488 – 498, 2004. URL `http://sag.sagepub.com/content/35/4/488.abstract`, last access 25/04/2011.

[206] Martha Koutri, Nikolaos Avouris, and Sophia Daskalaki. A survey on web usage mining techniques for web-based adaptive hypermedia systems. In S. Y. Chen and G. D. Magoulas, editors, *Adaptable and Adaptive Hypermedia Systems*, chapter 7, pages 125 – 149. IRM Press, London, UK, 2005.

[207] Domenico Talia, Paolo Trunfio, and Oreste Verta. Weka4WS: A WSRF-Enabled Weka Toolkit for Distributed Data Mining on Grids. In Alípio Jorge, Luís Torgo, Pavel Brazdil, Rui Camacho, and João Gama, editors, *Proceedings 9th European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD 2005)*, volume 3721 of *Lecture Notes in Computer Science*, pages 309 – 320. Springer Berlin / Heidelberg, Porto, Portugal, October 2005. 10.1007/11564126_32.

[208] ICEAGE Digital Library. EU ICEAGE project. URL `http://baillie.lib.ed.ac.uk:8080/`, last access 25/04/2011.

[209] Xen. URL `http://xen.sf.net/`, last access 25/04/2011.

[210] Stephen Childs, Brian Coghlan, John Walsh, David O'Callaghan, Geoff Quigley, and Eamonn Kenny. A Virtual TestGrid, or how to replicate a national Grid. In *Proceedings ExpGrid workshop at 15th IEEE International Symposium on High Performance Distributed Computing*, pages 47 – 54, Paris, June 2006. URL `https://mailserver.di.unipi.it/ricerca/proceedings/HPDC-15/EXPGRID/07Childs.pdf`, last access 28/04/2011.

[211] Stephen Childs, Brian Coghlan, David O'Callaghan, Geoff Quigley, and John Walsh. A single-computer Grid gateway using virtual machines. In *Proceedings of Advanced Information Networking and Applications (AINA'05)*, pages 310

– 315, Taiwan, March 2005. URL http://hdl.handle.net/2262/2398, last access 28/04/2011.

[212] Marcus Hardt and Rüdiger Berlich. Xen: Scientific Use Cases and Performance Comparisons. In *Proceedings UKUUG Linux Technical Conference*, Swansea, UK, August 2005. URL http://www.ep1.rub.de/~ruediger/xen-grid-in-a-box-ukuug.pdf, last access 28/04/2011.

[213] OpenCA. URL http://www.openca.org/, last access 25/04/2011.

[214] Owen Conlan. *The Multi-Model, Metadata driven approach to Personalised eLearning Services*. PhD thesis, Trinity College Dublin, 2004. URL http://www.scss.tcd.ie/Owen.Conlan/publications/Conlan_Thesis.pdf, last access 27/04/2011.

[215] Knowledge and Data Engineering Group (KDEG), 2008. URL http://kdeg.cs.tcd.ie/, last access 25/04/2011.

[216] Apache Tomcat. URL http://tomcat.apache.org/, last access 25/04/2011.

[217] Apache. URL http://www.apache.org/, last access 25/04/2011.

[218] GridSite, 2008. URL http://www.gridsite.org/, last access 25/04/2011.

[219] U.S. Department of Health and Human Services. *Research-Based Web Design and Usability Guidelines*. U.S. Government Printing Office, Washington, DC, 2006. URL http://www.usability.gov/guidelines/guidelines_book.pdf, last access 25/04/2011.

[220] Wolfgang Meier. eXist: An Open Source Native XML Database. In *Web-Services, and Database Systems, NODe 2002 Web and Database-Related Workshops*, pages 169 – 183, Erfurt, Germany, 2002. Springer.

[221] Quattor. URL http://quattor.web.cern.ch/quattor/, last access 25/04/2011.

[222] Lisa M. Lane. Insidious Pedagogy: How Course Management Systems Impact Teaching. *First Monday*, 14(10), October 2009. URL `http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/2530/2303`, last access 25/04/2011.

[223] HyperGraph. URL `http://hypergraph.sourceforge.net/`, last access 25/04/2011.

[224] Kathryn Cassidy, John Walsh, Brian Coghlan, and Declan Dagger. Using Hyperbolic Geometry for Visualisation of Concept Spaces for Adaptive e-Learning. In *Proceedings International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH2006)*, pages 421 – 426, Dublin, Ireland, 2006.

[225] SSLWithFormFallback Authenticator Valve for Tomcat. URL `http://wiki.apache.org/tomcat/SSLWithFormFallbackAuthenticator`, last access 20/02/2012.