DBTech Database Virtual Laboratory Workshops in Teaching*

Lili Aunimo, Martti Laiho, Aki Lassila
Business Information Technology Programme, HAAGA-HELIA University of Applied Sciences, Helsinki, Finland
lili.aunimo@haaga-helia.fi, martti.laiho@haaga-helia.fi, aki.lassila@haaga-helia.fi

ABSTRACT

In order to address the need to improve the quality and relevancy of studies in database technology, the DBTech Virtual Laboratory Workshops (VLWs) were developed. This paper describes the usage of these VLWs in teaching e.g. database administration and tuning, database modelling and semantics, concurrency control and recovery, business intelligence, and XML technologies. The benefits of using VLWs in teaching include e.g. their usefulness in demonstrating performance and scalability related issues from database administrator's and application developer's viewpoint.

Introduction

Databases are needed everywhere in the modern information society as they are the basis of all applications requiring reliable, secure and scalable data storage and access. The rapid change in the field of ICT technologies reflects in expanding topics in the software engineering curricula, while demand for optional subject areas and cross-discipline studies results in decreasing number of hours available for focusing on special studies such as database technologies. In addition, the academically oriented database textbooks typically used on the database courses tend to be quite theoretical, borrowing each other and repeating the applied mathematics such as relational algebra and classical serialisability theories (in fact born before the era of relational databases), of which the de facto database language SQL has out-grown with the procedural, object-relational, large objects, multimedia, and XML extensions, and the practical implementations of various concurrency control technologies reflecting the software engineering needs. The too theoretical textbooks with non-relevant materials in terms of use of the real DBMS systems used in the industry, together with the decreasing hype and attractiveness of the ICT in general in minds of young people, has resulted in decreasing interest and number of students – and therefore also in decreasing study hours allocated for database technologies. This is now leading to lack of competence in database technologies in the European industry. This calls for the need to organize database courses to reflect knowledge and skills on modern database technologies relevant to the needs in industry to our students in the limited number of hours allocated for these studies, without loosing the valuable time on non-relevant topics. The relevant and useful topics should be covered in attractive form to the students allowing experimenting, learning by doing, and learning by verifying.

The rapid change in the field of ICT technologies reflects also in the expanding field of database technologies, especially in the data access engineering. This has been a challenge to database professionals and also to teachers responsible of teaching the design and administration of databases, and especially the ever expanding field of various data access technologies in software development. In addition, typically there

* *This report is the result of collaborative work undertaken along the lines of the DBTechNet Consortium. All the co-authors participate in DBTech EXT, a project partially funded by the EU LLP Transversal Programme (Project Number: 143371-LLP-1-2008-1-FI-KA3-KA3MP)*

are only few teachers in a single educational institute working in these fields, leading to limited possibilities to collegial communication for teachers to keep in pace with the development of technologies. In 1996 we got worried on a non-sustainable development, as we noticed the expanding use of some non-mature data access technologies, especially among SME companies and students working in such companies. This led us, a group of database teachers, to search for co-operation between colleagues in European HE institutions in the LdV project called DBTech Survey, in order to keep with the pace of the development, to develop and share quality teaching materials, and to provide students and industry proper skills in robust use of the database technologies. As the future targets we defined "DBTNet", the network of teachers and professionals working in the field of databases, and "DBTLab", a set of shared laboratories making available among the partners the real database products used by the industry. The DBTNet was later renamed as DBTechNet in 2001 (see www.DBTechNet.org). In the subsequent LdV project DBTech Pro, as a form of developing shared teaching materials, we created the format of DBTech Pro workshops for testing and piloting the use of our tutorials and exercises. The first virtual laboratory over Internet was Index Design Laboratory, and as the pilot of the lab, the students at Malaga University experimented, over Internet, with some selected Oracle databases and index configurations located in Häme Polytechnic in Finland.

Since those days the PCs have become much more powerful and laptops are common among students, teachers, and private professionals. Current open source operating systems and free software available in Internet provide possibilities to build robust environments by everyone. Now even the mainstream database vendors provide free editions of their products, which can be installed also on laptops. Finally the advent of free virtual computer systems have made it possible for us to create virtual computer images, portable virtual laboratories with a free operating system, coupled with pre-installed databases and engineering tools, and share these over Internet without licensing problems. So now the decade old vision of DBTLabs has become true, and with co-operation we can fulfil the educational computing infrastructure requirements of "computing facilities at least equivalent to those used in industry" as stated in IS2002 Model Curriculum by Gorgone et al. (2002).

The current DBTech EXT project provides these tools for research, teaching, and studying in the form of virtual laboratory workshops (VLW, see e.g. Delgado et al. 2009) for higher level education and training on database technologies tuned to serve the needs of the database professionals and application developers. The VLWs educational and training content identifies, focuses on and addresses the essential knowledge areas of database technology, covering the related standards, specifications, and technology trends, thus promoting the understanding in practice of the mainstream DBMS technologies and products.

Related Work

Virtual Laboratories (VLabs) have been studied and used widely and for a long time in the framework of the education ICT professionals. This is in line with the recommendation of the Model Curricula for Computing (IEEE 2002) which states that practical laboratory activities are an essential part of any ICT curriculum. The term Virtual Laboratory has several definitions. Prieto-Blázquez et al. define a Virtual Laboratory as "an interactive virtual space that incorporates all the technological, pedagogic and human resources for carrying out practical activities, adapted to the

needs of the students and teachers in a virtual learning environment" (Prieto-Blazquez et al. 2008).

The Open University of Catalonia (UOC) has developed VLabs for its undergraduate degree programmes in Software Engineering, Computer Engineering and Telecommunications Engineering. Since the whole operation of the University is virtual, the need for VLabs is evident. The Virtual Programming Laboratories (VPLabs) offered by the university consist of the following parts: 1) technological resources, 2) pedagogic and strategic resources and 3) teacher resources. The technological resources consist, among others, of a remote laboratory that provides students with a real environment to carry out practical work. Remote laboratories are used when it cannot be assumed that students will have them at home or when it cannot be assumed that they can install them at home. Traditionally, remote laboratories have been developed for subjects such as computer networks and operating systems (Prieto-Blazquez et al. 2008).

The technological Resources of the VPLab of the UOC also consist of a virtual machine. A virtual machine is an environment that emulates the hardware of a complete physical computer. This enables the students and teachers to run multiple operating systems simultaneously on a single physical computer and to switch between them. Virtual machines also eliminate the need of performing practical activities such as software installation, software set-up and hardware set-up. Furthermore, virtual machines allow for the standardisation of the initial working environment. The virtual machine at UOC is implemented by installing the required software on a DVD that may be run without installation on a hard drive. This DVD is sent to the students and it contains free software such as the Ubuntu distribution of Linux and the Dev-C++ compiler.

In addition to the technical environment that contains the software to be learned and to be used in learning the theory behind the software, all e-learning environments have to take into account the need for interaction and collaboration. There are several ways in which this interaction can be supported by the learning system, and each system and each organization tends to have its own solution for this, see e.g. the virtual communication environment of the OUC (Prieto-Blasquez et al. 2008) and the web-based CSCL (computer-supported collaborative learning) environment named WeColLab (Bochicchio and Longo 2009). The tools used for interaction between students and between students and teachers typically contain both tools for non real-time (or asynchronous) interaction and for real-time (synchronous) interaction. Tools for non real-time interaction typically consist of tools such as:

- Assignment tool where a teacher may give assignments and the student may turn in assignments,
- Online discussion forum where students may post their questions and learning experiences, answer the other students' questions and discuss different topics,
- Online test tool where the teacher may create a test and the students may answer to the test,
- List of members, which typically contain the students and teachers of the course, their names, photos and other information. Their current status – online or not online – is typically also shown. This last feature could also be classified under the instant interaction tools.
- Collaborative editing environment such as a wiki, Google docs, spreadsheets etc.

The real-time interaction tools typically contain tools such as:

- Instant messaging tools. Widely available free software such as Google Talk may be used for this or LMS system specific chat tools may be used.
- An editing tool that supports synchronous editing, such as the web-based freely available word processor Etherpad
- Video chat and conference call tools such as Skype and Adobe Connect Pro.

In this paper, we use e-learning as our conceptual framework as we study how the DBTech VLWs can be applied in university level teaching and in offering (online) courses in higher education. We present a view that we need to take into account the different stakeholders' views and their relationships in order to create successful VLWs that can be used in teaching meaningfully. We consider three different stakeholders that are the universities' faculties, developers of the VLWs, and, naturally, the students. From their perspectives, the main questions that need to be addressed are how to use VLWs meaningfully in (especially online) teaching, how to develop educational VLWs, and how to utilise these VLWs in enhancing learning.

We see e-learning as an ICT-enabled form of education utilising the Internet as a tool and channel for providing tools of learning and teaching. Roots of e-learning date back to the use of Computer-Based Training (CBT) that utilised specific training programs in classrooms. As the Internet technologies evolved, also e-learning began to develop as a means of distributing and sharing of information and instruction to a wider audience. The learning environments have developed, but usually the teachers are still lacking in understanding how to apply e-learning in practice and in this case with VLWs. Based on this, it can be said that the use of e-learning is still in a stage of exploration (Britain and Lieber 2004) and therefore a further study on the usage of the VLWs in teaching is in order.

This case study follows the guidelines of interpretive approach and e-learning as an educational paradigm is the conceptual framework through which the findings are analysed by using hermeneutics as our mode of analysis (Klein and Myers 1999). The unit of analysis in this case study (Yin 2003) is the DBTech VLW. Based on our observations/empirical data we try to answer the research question: what are the benefits and challenges for the stakeholders related to using VLWs in their teaching and learning?

Virtual Laboratory Workshops

Virtual laboratory workshops (VLW) developed in the DBTech EXT -project are learning environments consisting of a virtual machine and of an e-learning environment. The virtual machine typically contains an operation system (Linux), database software (e.g. DB2 Express-C), possibly an accompanying application such as data mining software (e.g. WEKA), a set of software that can be used to populate the database with data or to run tests and possibly a set of sample data. The VLW package may also contain a tutorial for the study subject. As standalone test beds, VLWs are also suitable for self-study in online or offline fashion.

In addition to a virtual machine, many of the VLWs also contain an e-learning environment such as an environment created in the Moodle Learning Management System (LMS). This is the case in e.g. the VLW on Business Intelligence (Dervos et al. 2009). Moodle is an open-source LMS that requires the PHP programming language and a database to be installed in the server on which it is installed. It has many modules and plug-ins that can be used to customize it. Modules and plug-ins may naturally also be developed, if needed. The simplicity of installation, the active developer and user communities as well as the wide variety of available plug-ins and

the possibility to develop new plug-ins make Moodle an attractive LMS to be used as an environment in teaching database-related subjects. Data mining on Moodle data, such as student's usage data, is also possible (Romero et al. 2007). This makes Moodle an attractive environment for the teacher who is interested in developing his educational practices and in contributing to the research on educational usage of e-learning environments.

In VLWs organized in the Moodle environment, we offer the following tools for asynchronous communication: a newsgroup, list of course participants and teachers with a photo and other additional information, course assignments with a digital drop box and a way of giving feedback to the student as well as the course tutorial materials. As a means of synchronous communication, we offer prescheduled chat sessions. The DBTech VLWs may use the instance of the Moodle software that is installed as the DBTech Portal server at the University of Macedonia, or the instances of the Moodle software that are installed on the servers of the participating organizations. This is possible because Moodle supports the export from one Moodle installation and the import to another Moodle installation. This may be achieved for example by using the IMS Content Packaging Standard to package the learning environment. Using standard packaging formats also enables the re-use of the contents in other learning environments such as Blackboard and Sakai.

Experiences

Virtual computers provide many opportunities as learning environments on different courses of database technologies, such as system and database administration, database design, and/or data access engineering as follows: A new virtual computer without even operating system can be used for exercises on installing the operational environment and networking the computer. A virtual computer with the basic infrastructure of operating system and basic tools can be used for exercises on installing a database server with administration client, database modelling and development client middleware and application development tools. Finally a ready installed database server with databases, tutorials, and exercises can be used to practice use of the native SQL language of selected DBMS and various data access technologies. The same can be applied both on courses of basic OnLine Transaction Processing (OLTP) databases, and Data Warehousing (DW) and Business Intelligence (BI) including OLAP, Data Mining and Knowledge Discovering from Databases (KDD). In this presentation we will focus on needs of basic OLTP courses.

In DBTech EXT project we have used two kind of sharing of virtual computers: remote servers, and downloadable virtual images. University of Malaga (UMA) has provided us opportunity to create virtual computer images on their Virtual Desktop Infrastructure (VDI) server, which images can be accessed by registered remote users who need only browser access to the server. This has been described in the paper of Dervos et al (2009). This approach allows use of licensed software by limited participants, and applies only e-Learning courses arranged by the HE partners on specific academic programmes, such as Academic Initiative of IBM.

A mission of DBTech EXT is to prepare free learning materials for students and database professionals to be used on their lifelong learning. For this purpose we have been experimenting with the following free virtual computer software: VMware Player of VMWARE Inc., Virtual PC 2007 of Microsoft, and VirtualBox of Sun. In these virtual computers we have created free virtual laboratories using only free software: Linux variants as operating systems, the latest free editions of DB2 and Oracle DBMS systems, Java developments software, and tutorial materials of our

own. These virtual laboratories we have put on our web site in order to be freely available at www.DBTechNet.org.

The VLWs of DBTech EXT are based on live workshops organized first at some partner HE institute of DBTechNet. Then the workshop material has been rewritten into self-study material, which can be used either in classroom environment on courses, or which can be downloaded and installed on private computers of single, independent learners, such as adult database professionals. For the VLWs we have written tutorials on database administration and data access technologies of OLTP databases. These tutorials can be used in the ready downloadable virtual laboratories as a private VLW. These downloadable free VLWs, using only free software, include currently the following: Database Recovery (tutorial), Index Design (SUSE VLab + tutorial + DB2 exercises), SQL Concurrency Technologies (Debian VLab + tutorial + exercises), Data Access Patterns (Debian VLab + RVV Paper + slides), TPCA2 Benchmark (Debian VLab + tutorial + instructions), ORM (Debian VLab + tutorial + exercises), and XML and Databases (Debian VLab + tutorial + DB2 exercises).

Case: VLWs and TPCA2 Benchmark

As a contents example of our VLWs we will explain the case of the TPCA2 Benchmark, which is based on the TPC-A benchmark by Transaction Processing Performance Council (see www.tpc.org). The idea of a standard database transaction benchmark for comparison of throughput and price/performance on different database environments on selecting solution for an OLTP application including hardware, operating platform and a DBMS was first presented by Gray (1985). Based on the article, the official TPC benchmark TPC-A was introduced in 1989. TPC-A uses a simplified transaction in a bank account system. The specified transaction is repeated by a number of concurrent client sessions, and performance is measured in metrics like local and global Transactions Per Second (TPS) units, and also in cost per transaction. However, already in 1995 TPC-A was declared obsolete since more advanced benchmarks were developed by the TPC council.

Orfali and Harkey (1998) presented a simplified version of the TPC-A written in Java and using SQL Server as the DBMS. Instead of benchmarking DBMS systems, they presented their version as Debit-Credit Benchmark tool for experimenting differences in performance when transaction is based on using SQL commands or implemented as a store procedure to minimize the round-trips, including the network traffic between the client and database server. The simplified Debit-Credit benchmark uses a scalable BANK database of 4 tables. The smallest scale "1-tps" is presented in Figure 1 below, having 10 branches, 100 tellers, and 100 000 accounts.

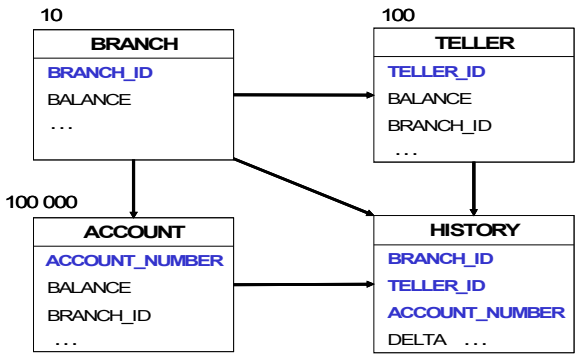

Figure 1: Tables of the TPC-A BANK database of scale 1-tps

The TPC-A transaction selects a random row from the Account table and a random row from the Teller table. The account data defines the row of its home branch office in Branch table. Then balances of the account, the teller, and the branch rows are updated, and information of this bank transaction is written in the History table.

We saw the potential of Debit-Credit as a learning tool to experiment with data access varying connection strategies, isolation levels, and behaviour of different DBMS systems. Therefore, we adapted the Debit-Credit benchmark as 2-Tier client/server and 3-Tier J2EE as well as Microsoft MTS application server implementations using these as learning tools in database programming courses at Helia (now Haaga-Helia), calling our modified version as TPCA. This was presented in 2003 at T.E.I. of Thessaloniki (now ATEI) in the first live workshop of DBTech Pro, where the participants, synchronized according to instructions of the workshop instructor, started sequences of concurrent transactions in a DBMS server using the same strategy of transaction parameters, observing performance in TPS units, number of concurrency conflicts (deadlocks) and number of retries needed to get every transaction succeeded. This was an exciting experience to students and partners in the workshop, and since then it has been repeated by DBTechNet partners as an exercise on transaction processing performance.

The problem of TPCA benchmark was that it required a group of users running the client software concurrently, so it does not suit to be used as a self-study exercise in a virtual laboratory. Therefore we have now modified the lab into TPCA2 benchmark, in which the learner defines parameter values for the transaction strategy to be applied in the test, including series of concurrent transactions at a time, and starts the experiment. The client program follows the parameter settings creating a number of concurrent threads accessing the database and collects statistics of the experiment in the same or a different database.

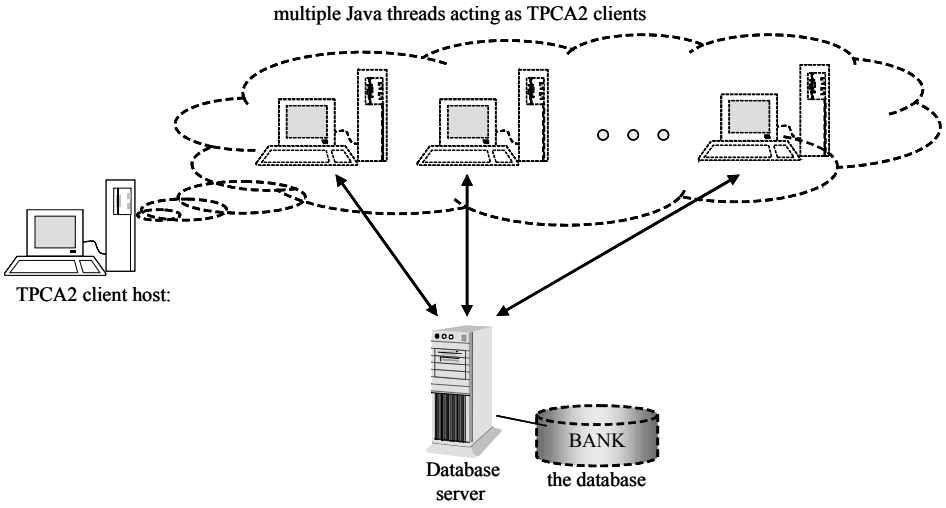

Figure 2: The new version of TPCA2 benchmark test (included in the VLWs)

After the test the statistics are collected from the database and presented in a chart form on a worksheet. A sample experiment on DB2 Express-C run on VirtualBox Debian Linux is listed Figure 3. The client host written in Java creates in this experiment 2-10 threads at a time. Each thread uses one the following 6 strategies on running the TPCA transaction 40 times using the following transaction

implementation alternatives: 1) direct SQL Statements, 2) pre-prepared statements with execution plans in the cache of the DBMS, 3) calling the stored procedure of the transaction.

Furthermore, in case of "single connection" the thread repeats the transaction using a single database connection, which simulates use of connection pool, whereas as "connection/transaction" means that a new connection is opened for every run of the transaction.

	TPS statistics		DB2 Express-C 9.7		VirtualBox Debian	
	Statements		PreparedStatements		Stored Procs	
	Conn/thread	Conn/trans	Conn/thread	Conn/trans	Conn/thread	Conn/trans
Threads	SQL_c1	SQL_c/Tr	pSQL_c1	pSQL_c/Tr	proc_c1	proc_c/Tr
1						
2	12,0	8,2	34,0	17,1	63,0	20,6
3	8,9	6,2	32,3	12,0	46,3	14,6
4	6,9	5,2	24,6	9,5	42,5	12,2
5	5,4	4,2	21,6	7,6	32,3	9,7
6	5,2	3,3	19,3	5,7	36,9	8,1
7	4,4	3,2	18,2	5,7	36,6	6,8
8	3,7	2,7	17,7	5,1	29,9	6,9
9	3,5	2,4	16,9	4,7	31,0	6,2
10	3,2	2,2	16,4	4,2	29,9	5,5

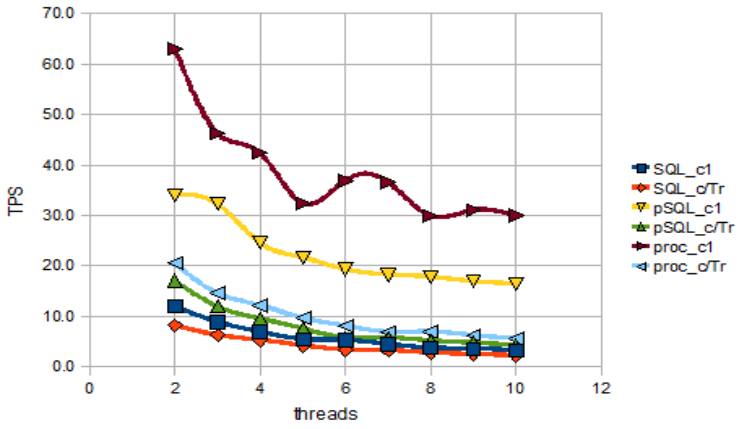

Figure 3: Example of the performance and scalability results

The TPS values measure performance of the system in terms of throughput depending on the selected transaction strategy and isolation levels for number of concurrent transactions (i.e. number of threads). The better the performance stays while number of concurrent transactions increases the better scalability the system has.

It has to be noted that the results vary between experiment runs, but they provide learners the overall idea of the effects of the various tested performance factors. By experimenting with the different DBMS systems the learners will find out that the DBMS systems behave differently. The results may be surprising even to database professionals. However, we should not compare between the DBMS systems based on our virtual laboratory experiments. These servers may not be tuned for their best performance. Also if we run the experiments in a single virtual computer, the effects of disc I/O and network traffic are not real. Remembering these facts, the results are comparable only between various experiment configuration runs in the same DBMS system, in the same virtual computer, and in the same physical host computer.

Implications

By using the virtual laboratory workshops the students can study the different factors in data accessing. For example the effects on performance and scalability when using connection pooling, stored procedures, command caching, and transactions. The benefits of using VLWs are that they are fully controllable environments for conducting benchmark tests, i.e. they contain a database, data, and the variables that are used in order to illustrate the effects that e.g. using stored procedures have on the performance of the database itself.

The first version of TPCA2 benchmark has been in use at Haaga-Helia on Data Access Technologies course and by some students on their BSc. studies experimenting with DB2, Oracle, SQL Server, MySQL, and Derby DB. Own observations clearly make the results more interesting to students. Some results make students curious of the reasons and lead to extended studies of the phenomena behind the results. There is more to be discovered from the collected statistics and we need to extend tracing and reports of the experiments in the future. We expect that the benchmark and instructions will get better as we learn of the use by more new learners and teachers.

Considering environmental sustainability, virtual computers help us to preserve environment from the ever increasing number of used computers, since on a terabyte portable USB disc one can have tens of different virtual computers, which can be easily ported between home, school, laboratories, etc. Instead of filling a room with different servers, whole network of servers may be running on a reasonably equipped workstation.

Discussion

Earlier, before the development of these DBTech VLWs, the illustration of these database management and application development related issues that were taught in educational institutes in theoretical level only and could not be done in practice. In addition, many application developers and teachers of database technology did not have exact knowledge of the magnitude of the effects of these different factors.

Students of today have grown accustomed to action and learning by doing. The virtual laboratories support this experimental style of learning. The study problems can be seen as games which can be solved in the virtual laboratory using real DBMS systems of industry, instead of simulations or just plain theories and textbook exercises. Also a virtual laboratory allows studying the behaviour, administration, and use of the DBMS systems with full administrator privileges, which is not possible in secured industry environments, whereas virtual computers provide safe learning “sandboxes”, attractive also to professionals. If something breaks in the system, the whole virtual computer can be purged, and the study work can continue with a fresh new copy of the virtual computer.

From the participating universities’ viewpoint, they have succeeded in integrating VLWs as tools for teaching database management related functionalities and issues. Additionally, DBTech VLWs are also viewed as good tools for providing online courses as they can be used also outside the campus areas for self-study purposes. However, benefiting more efficiently from the use of VLWs in e-learning will require more experiences and study in order to provide teachers both a deeper understanding of e-learning pedagogy and the usage of DBtech VLWs in their own teaching. As the scope of this case study was limited, these findings are initial and further studies in other educational organisations beside the participating universities would be needed in order to produce more generalisable results applicable in practice.

Conclusions

This paper presents that by using the DBTech virtual laboratory workshops, the students now have the means of gaining practical, hands-on knowledge of how the different methods of accessing and using database affect performance and scalability. Furthermore, these VLWs help to illustrate the theoretical issues of data access strategies in building applications that use databases.

The purpose of this interpretative study was to study DBTech Virtual Laboratory Workshops in order to identify the different stakeholders' benefits and challenges associated with using WLVs in teaching and how they can be applied in the e-learning context. Our initial findings show that universities have benefited from using the VLWs as they have enabled them to provide e-learning courses related to database management functionality to students efficiently and in more practical level. Also, using the VLWs seems to enhance the pedagogical quality of their (online) teaching. However, in order to fully exploit the advantages of e-learning the focus should be shifted more from technology to teaching. In other words, creating these VLWs is not enough but understanding the pedagogic principles of e-learning and teaching using the VLWs should be the focus. It should also be kept in mind that e-learning is still in its development phase and also (unfortunately) not that many members of the universities' faculties can use the LMSs well nor provide e-learning in meaningful way because they lack the training and necessary skills. Therefore, we recommend that further studies should be conducted in order to study how DBTech Virtual Laboratory Workshops should be used in providing (online) DB management courses.

References

Britain, S. and Liber, O. (2004) *A Framework for the Pedagogical Evaluation of Virtual Learning Environments. JISC Final Report*. (www.jisc.ac.uk/uploaded_documents/jtap-041.doc)

Delgado, I., del Mar Roldan Garcia, M., Farfan-Leiva, J., Laiho, M., Laux, F., Dervos, D. and Aldana Montes, J. (2009) *Innovative Unity in the Diversity of Information Management Skills teaching and Training across Europe*. IADAT Journal of Advanced Technology on Education, 3(3), 439-434.

Dervos D.A., Evangelidis G., Karamitopoulos L., Aunimo L., Aldana-Montes J.F., Farfán-Leiva J.J., Siakas K., Valkama H. (2009), DBTech EXT Virtual Lab Workshop: Business Intelligence, *Proceedings of the eRA-4 International Conference on the Contribution of Information Technology to Science, Technology and Education*, September 24-26, Spetses, Greece.

Gorgone, J., Davis, G., Valacich, J., Topi, H., Feinstein, D. and Longenecker, H. (2002) *IS 2002: Model Curriculum and Guidelines for Undergraduate Degree Programs in Information Systems*. Association for Information Systems, (http://www.acm.org/education/is2002.pdf)

Bochicchio, M. and Longo, A. (2009), A Multi-purpose Architecture for Collaborative Web Labs, *Proceedings of Ninth IEEE International Conference on Advanced Learning Technologies* (ICALT), July 15-17, Riga, Latvia.

Prieto-Blazquez, J., Herrera-Joancomarti, J. and Guerrero-Roldán, A.-E. (2009) *A Virtual Laboratory Structure for Developing Programming Labs*, International Journal of Emerging Technologies in Learning, Vol. 4 (2009).

IEEE Computer Society / ACM Task Force on Computing Curricula (CC). *Model Curricula for Computing*, IEEE Computing Society Press and ACM Press, 2005.

Klein, H. and Myers, M. (1999) *A Set of Principles for Conducting and Evaluating Interpretive Field Studies in Information Systems*. MIS Quarterly, 23(1), 67-94.

Romero, C., Ventura, S. and García, E. (2008), *Data mining in course management systems: Moodle case study and tutorial*. Computers & Education, 51(1), 368-384.

Yin, R. (2003) *Case Study Research: Design and Methods*. Sage Publications, Thousand Oaks, CA.

Orfali R. and Harkey D. (1998) *Client/Server Programming with JAVA and CORBA*, 2. edition, John Wiley & Sons Inc., New York, NY.

Gray, J. (1985) *A Measure of Transaction Processing Power*, Datamation, 31(7), 112-118.