Learning GRID

A newsletter form the Kaleidoscope Learning GRID SIG **Chief Editor: Nicola Capuano**

Newsletter #6: October 2005

Deliverable D 12.2.8

Prepared for the European Commission under Contract No. NoE IST-507838 as a deliverable from

Task: 12.2: Networking and Dissemination
Date of issue: 31/10/2005
Version: 1.0
Distribution: public

In this Issue:

Editorial

Welcome to the sixth issue of the Kaleidoscope Learning GRID SIG newsletter.

Now, that the summer is over and that all the students have returned to their lessons, we are also back to publish this issue which will be the last for this year. As we are reaching the end of 2005, we already have one year and a half experiences to share.

As it has been mentioned in previous issues, the learning process is quite complex. It has been proved that sharing experiences and thoughts among students help them to understand some concepts.

The activity led in this SIG GRID within the Kaleidoscope Network of Excellence is a clear example for this; the more we share with the others the more we learn. Therefore the aim of this newsletter is to share with all our readers the knowledge and experiences that we have acquired about GRID applications and services for e-learning.

In particular, in this issue we would like to share with you an article from one of the groups that compose the SIG GRID (the University of Catalonia, Spain). In its article, this group presents a parallel Grid-based implementation for processing in real time the event log data generated in collaborative applications. They also present the results when applying this to real scenarios with relevant conclusions that show the feasibility of using grid middleware to speed and scale up the process of structuring and processing semi-structured event log data.

For our section "Research project focus" we have chosen the GridCole project that shows us new scenarios and applications for collaborative learning. This is possible thanks to new tailorable applications which are being tested at University of Valladolid, that can guide students to achieve a common final result, allowing them to collaborate among different groups.

In the "Technological View" section we present the Globus Project and the Globus Toolkit. This toolkit is considered as the de facto standard for the development of Grid services. We think that most of you who are dealing with Grid applications know it, however here we relate a bit of its history and present the new features offered by the last release.

Finally, and to keep you informed and updated, we include in our usual sections of "News" and "When What Where" the latest pieces of information relevant for your research group and the events that you may consider in your agenda.

In the vicinity of the Online Educa Berlin event (which is nowadays one of the compulsory dates for those dealing with education and training, especially within the industry) we hope that you will find this issue interesting and plenty of good materials to share with others, enhancing therefore the learning process for all.

Enjoy your read.

Blanca Jordan

Learning SIG Grid Member

Grid-based Implementation for Processing Data in Collaborative Applications

The goal of this article is to present a parallel Grid-based implementation for processing in real time the event log data generated in collaborative applications. Our results show the feasibility of using grid middleware to speed and scale up the process of structuring and processing semi-structured event log data.

1. Introduction

In Computer-Supported Collaborative Learning (CSCL) environments, the analysis of the information related to the collaborative group activity is crucial for understanding collaboration and group processes [1]. The Basic Support for Collaborative Work system [2] is a collaborative application that generates event log data regarding the connection information as well as actions performed by the users during a certain period of time.

The information generated by the collaborative applications can be of a great variety of type and formats [3]. Moreover, collaborative applications are characterized by a high degree of user-user and user-system interaction and hence generate a huge amount of information of event log data.

As a matter of fact, the computational cost is the main obstacle to processing this data in real time [4] and in real situations this processing tends to be done offline in order to avoid harming the performance of the logging application, but as it takes place after the completion of the collaborative activity has less impact on it [5]. Most of the existing approaches in the literature consider a sequential approach for the processing of log data and try to overcome the performance problem by: (i) processing for specific purpose (i.e. limiting the quantity of information needed for that purpose); (ii) processing of small data samples, usually for research and testing purposes.

The concept of a computational Grid [6] has emerged as a way of capturing the vision of a network computing system that provides broad access to massive computational resources. Thus, in this paper, we show how easily we were able to offload onto the grid the online processing of log data from the collaborative application and how a simple Master-Worker scheme sufficed to achieve considerable speed-up. In order to show the feasibility of our approach, we use the event log data from BSCW system in our real context of Open University of Catalonia [7] though our approach is generic and can be applied for structuring event log data of collaborative applications in general.

2. The problem of structuring and processing event log data

The problem of structuring event log data of collaborative applications in real time can be defined as follows: give structure to the semi-structured textual event log data that an application logs as soon as it logs it and persist the resulting data structure for the later processing by analysis tools.

A special case of this problem, known as log data normalization or unification that consists in transforming proprietarily formatted log data to a standard log data format, is recently gaining attention from the autonomic computing community [8] as a way to give standard and homogeneous structure to the strongly heterogeneous data that the disparate elements of an IT infrastructure log while they operate.

In fact in the discussion that follows we will be using the terminology and architecture of the Generic Log Adapter (GLA) [9], a framework that addresses the problem of real time log data normalization and structuring in the IBM's autonomic computing toolkit [10] for building autonomic systems and that has been open sourced as part of the Eclipse Hyades project [11]. The GLA is architected around four components mapping to each one of the four phases involved in the structuring of log data in real time: a sensor, an extractor, a parser and an outputter. More formally, the input of the problem of structuring log data is text, thus it can be modeled using formal language terminology [12].

Let the input be represented by a word, ω, from a given alphabet, Σ. The Sensor component reads this word as it is being generated, thus outputs a sequence of subwords of ω, say, $\omega 1$, $\omega 2$, …, ωm . The Extractor component acts on each one of these subwords one at a time, outputting a subword, $E(\omega i)$,of ωi which verifies one simple but important property: it is an independent unit of structure. That is, it contains all the information the Parser component needs to know in order to be able to transform it into a data structure. Hence the Parser component acts on it outputting a data structure, $P(E(\omega i))$, to the Outputter Component who persists it.

It's worth noticing here that the time complexity of the computation of $E(\omega i)$ is linear. Indeed, this computation is a word recognition problem. Thus, the question of its time complexity is reduced to what kind of languages might an Extractor Component ever had to recognize. It can be argued, due to the nature of log data, that these can only be Regular languages or, in the case of log data with multiple formats, a union of Regular Languages, which is also a Regular Language. Hence time complexity is linear at most [13]

3. A sequential approach

In order to deal with the problem of extracting useful information from the event logs generated by the BSCW system in real online learning group activities conducted at the Open University of Catalonia we have developed a simple application in Java, called EventExtractor.

This application runs offline on the same machine as the BSCW server and uses the daily log files generated by the BSCW server as input so as to: (i) identify the event boundaries inside the log file (extractor component), (ii) map specific information contained in these events about users, objects, sessions, etc. to typed data structures (parser component), and (iii) store these data structures in a persistent support (outputter component). Note that as the processing is done offline there is no need for a sensor component.

In order to analyze the performance of this sequential application and compare it with its parallel Grid-based (see Sect. 4), we designed a specific test battery in which we used both large amounts of event information and well-stratified short samples consisting of all the existing daily log files making up the whole group activity generated during the spring term of 2004 in the computer science subject "Software Development Techniques" at the Open University of Catalonia. This course involved two classrooms, with a total of 140 students arranged in groups of 5 students and 2 tutors. On the other hand, other tests involved a few log files with selected file size and event complexity forming a sample of each representative stratum. This allowed us to obtain reliable statistical results using an input data size easy to use.

All our test battery was processed by the EventExtractor application executed on single-processor machines involving usual configurations. The battery test was executed several times with different work load in order to have more reliable results in statistical terms involving file size, number of events processed and execution time along with other basic statistics. The experimental results from the sequential processing of eight event log files are summarized in Figure 1, where for each event log file we show the relative comparison scale for the file size, number of events and the processing time.

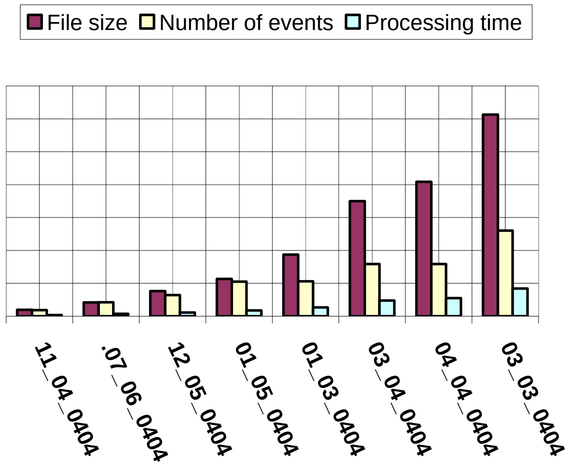

Figure 1: Sequential processing of event log files(file size in bytes, time in seconds)

In this figure a certain degree of linearity can be noticed between the number of events and the file size with regard to the processing time.

In a similar way, Figure 2 presents the processing results of over one hundred event log files involving file size and processing time showing that the processing time is linear on the size of the log file processed. This allows us to talk about the processing rate, P, (i.e. in Kb/sec), of the EventExtractor application, and to express the running time (i.e. in seconds) of the EventExtractor application with the following formula:

$$TS(n) = n/P$$

where ***n*** is the size (in Kbs) of the event log file.

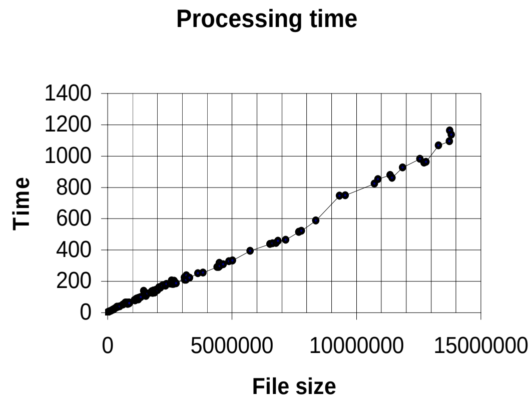

Figure 2: Sequential processing time (in seconds) vs. event log file size (in bytes)

4. A Parallel approach using the Master-Worker paradigm

In this section we present a high level overview of how the problem of structuring the event log data can be parallelized using the Master-Worker (MW) paradigm [14].

The MW paradigm has been widely used for developing parallel applications. In this model there are two different types of entities: master and worker. The master decomposes the main task into subtasks (sometimes this reduces to splitting the problem's input into parts) and sends these to the workers. The workers process the subtasks as soon as they receive them and send back the result to the master, which uses them in its main flow of computation.

The MW model has proved to be efficient in developing parallel applications with different degrees of granularity and is particularly useful when the partitioning of the problem is easy to compute and the dependencies between tasks are low. Indeed, this is the case of the problem of structuring plain text log data since (i) the Extractor Component outputs independent units of structure (i.e. messages/events) which means that if the problem is partitioned using the boundaries of these units no dependencies between tasks will exist, and (ii) the problem's input can be easily partitioned in these units of structure since, as we have seen, this can be done using regular expressions.

Given all the above, the problem of structuring log data in real time can be naturally parallelized using the MW paradigm by grouping the Sensor and Extractor components at the Master side and leaving the Parser and Outputter components at the Workers side.

One drawback of this approach, however, is that the Master is not in full control of the size of the task that it sends to workers since events/messages can have arbitrary size. This fact can somewhat reduce the capacity of the Master to play with the task size in order to better adapt to the different computational capacities of the workers and/or to their variable workloads, specially if the sizes of events/messages are too large. However, the later is not a property we could normally expect of log data.

5. A Parallel Grid-based approach

In order to experimentally test the feasibility of the MW paradigm for parallelizing the structuring of event log data we have implemented a minimal Grid prototype that parallelizes the EventExtractor application (see Sect. 3). We used the Globus Toolkit 3.2 and we deployed the prototype on the Planetlab platform. Both Globus-Toolkit 3.2 and Planetlab are briefly described next.

The Globus Toolkit (GT) [15] is the actual de-facto Grid middleware standard. Version 3 of GT (GT3) is a refactoring of version 2 in which every functionality is exposed to the world via a Grid service. Grid services are basically stateful web services. The core of the GT is a Grid service container implemented in Java that leverages and extends the Apache's AXIS web services engine.

Planetlab [16] is an open platform for developing, deploying and accessing planetary-scale services. It is, at the time of this writing, composed up of 528 nodes hosted in 249 different sites. Each Planetlab node is an Intel IA32 machine that must comply with minimum hardware requirements (i.e. 1GHz PIII + 1Gb RAM) running the same base software, basically a modified linux operating system offering services to create virtual isolated partitions in the node, called slivers, which look to users as the real machine. Planetlab allows every user to dynamically create up to one sliver in every node, the set of slivers assigned to a user form what is called a slice. It is said that a Planetlab node can run up to 100 concurrent slivers.

In order to test our Grid prototype we turned Planetlab into a Grid by installing the GT3's Grid service container in every sliver of our slice. Moreover, we implemented the worker as a simple Grid service playing the role of the parser and outputter components and deployed it on the GT3's container of every sliver of our slice. On the other hand, we wrote a simple Java client playing the role of the master and mapping to the sensor and extractor components which dispatches, using a simple list scheduling strategy, the tasks to the workers by calling the operations exposed by the worker Grid services.

Notice that our objective was not to create a full-blown GT3 MW implementation but rather to show the feasibility of a parallel Grid-based implementation using the MW paradigm for our problem domain, as follows.

The worker Grid service publishes an interface with only one operation, processEvents. The master calls this operation in order to dispatch a task to the worker. The worker can only do one of these operations at a time (no multithreading). The operation has only one argument: a string containing the textual representation of the events to be processed by that task. The operation returns a data structure containing performance information about the task executed (elapsed time in ms, number of events processed and number of bytes processed). The processEvents operation is implemented by wrapping the Java code of the EventExtractor application's routine that parses the BSCW log events. In other words, the workers execute exactly the same java bytecode to process the log events as the EventExtractor application. This makes possible the performance comparison between the sequential and Grid approaches.

On the other hand, the master is just a "normal" Java application that reads from a configuration file: (1) the folder that contains the event log files to process; (2) the available workers; (3) the number of workers to use, and (4) the size of the task to be dispatched to each worker expressed in number of events. The master then proceeds as follows: peeks as much workers as needed from the configuration file and puts them all in a queue of idle workers, then enters a loop reading line by line (i.e. sensor component) the data contained in the event log files located in the folder specified in the configuration file, and parsing each one of these lines in search of the boundaries between events in order to extract those (i.e. extractor component). Every time the master reads a number of events equal to the size of the task specified, it creates a thread that gets a worker from the queue of idle workers (synchronously waiting for a worker if the queue is empty) and synchronously calls the worker's processEvent operation. Once the call to the worker returns, the worker is put back into the queue of idle workers. The master exits the loop when all events in the event log files have been read and all the tasks that were dispatched are completed.

Notice that the scheduling strategy (i.e. list scheduling) favors the faster nodes and thus it is appropriate for an environment where worker machines have unpredictable workloads as the Grid, however, in a more homogeneous workload environment a simple static round robin scheduling strategy could be more efficient.

6. Experimental Results

In this section we present the experimental results of our Grid prototype, in order to evaluate them it is important to understand how they were collected and what was measured.

Basically we measured parallel speed up and efficiency for different executions of the parallel processing of 1000 events using different number of workers, p, (physical scaling) ranging in {2,4,8,16} and different task sizes, ns (i.e. in number of events) ranging from 1 event to $\lceil 1000/p \rceil$ events.

Parallel speed up is used to measure the performance gain from a parallelized execution of the application over its serial execution, defined as follows:

$$S(n,p,s) = T_S(n) \,/\, T_P(n,p,s),$$

where n is the size of the input, s is the task size, $T_S(n)$ is the total running time of the sequential execution for an input of size n and $T_P(n,p,s)$ is the total running time of the parallel execution for an input of size n, using p workers with a task size of s.

Parallel efficiency measures the degree of utilization of the computing resources involved in the parallel computation and is defined as the speed up divided by the number of computing resources (i.e. workers):

$$E(n,p,s) = S(n,s) \,/\, p.$$

To characterize the speed up of our prototype we run k different executions for each combination of number of workers and task sizes that we tested and then applied the following formula:

$$S_o(n,p,n_s) = \Sigma^k_{i=1}\Sigma^p_{j=1}T_{pEij}(n) \,/\, \Sigma^k_{i=1}T_{MWi}(n,p,n_s),$$

where $T_{pEij}(n)$ is the time spent by the j^{th} worker executing its *processEvent* operation in the i^{th} execution, while $T_{MWi}(n,p,n_s)$ is the total running time of the master in the i^{th} execution.

Note that we use averaged values to compute the speed up and that the total sequential execution time was not computed by running the *EventExtractor* application in an arbitrary machine, but by summing up the times spent by the workers executing its *processEvent* operation. This time can be thought of as the serial execution time of the application on a hypothetical machine with varying computational power and workload equivalent to the ones experimented by the workers during the parallel execution. Thus, we can have a more realistic idea of the speed up achieved.

We resume in Figures 3 and 4.x below the main results of our experiment. Figure 3 shows how the observed speed up and parallel efficiency of our prototype scaled with the number of workers for a fixed task size of 25 events, while Figures 4.x show, for each number of workers we tested, how the observed speed up varied with the size of the task.

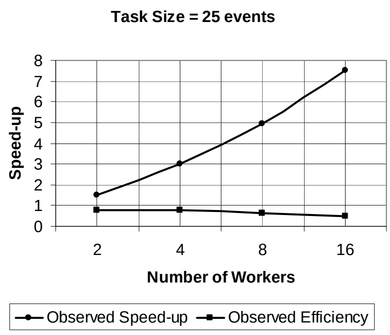

Figure 3: Speed-up and efficiency vs. #workers for a task size of 25 events

From the results obtained it can be concluded that a reasonable speed up has been achieved in every tested configuration. However we observe that the parallel efficiency decreases with the number of workers. This could be explained due to the fixed size of the input to 1000 events since the speed up seems to grow with the task size except for values near $\lceil 1000/\boldsymbol{p} \rceil$ where it begins to decrease.

Indeed, for too small values of the task size the overhead introduced by the transmission protocol when sending the parts to the workers is noticeable and the implemented list scheduling strategy may be spending too much time waiting for completion notifications. On the other side, values of the task size close to $\lceil 1000/\boldsymbol{p} \rceil$ considerably diminish the attainable degree of concurrency; however, it is here where increasing the size of the problem, ***n***, could be useful.

We finally note that although the results of this experiment are dependent on the form of the BSCW event log files, the parallelization strategy presented in this paper is generic and can be applied to parallelize the structuring of collaborative application's events log data.

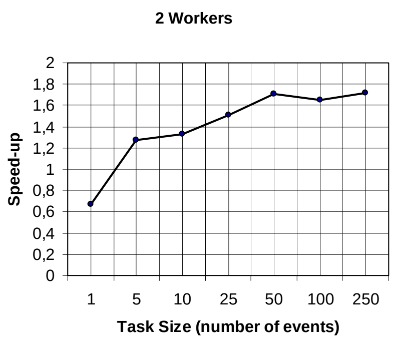

Figure 4.1: Speed-up vs. task size for 2 workers

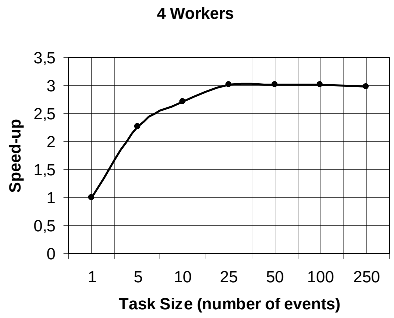

Figure 4.2: Speed-up vs. task size for 4 workers

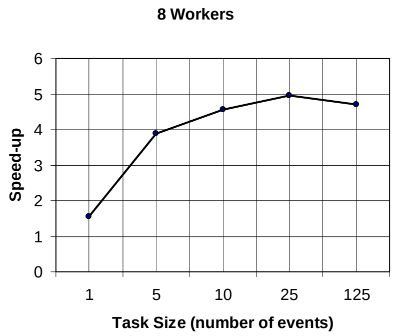

Figure 4.3: Speed-up vs. task size for 8 workers

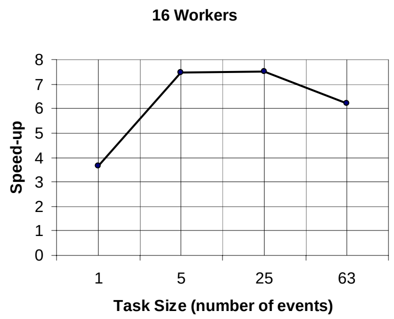

Figure 4.4: Speed-up vs. task size for 16 workers

7. Conclusions

Our results show the feasibility of parallelizing the problem of structuring any plain text event log data, achieving considerable speed up. Aspects of the prototype that we plan to enhance are fault-tolerance, dynamic discovery of workers and the possibility of implementing the communication between the master and the workers by other means than the default transport mechanism (i.e. SOAP over HTTP) used by GT3.

References

1. Dillenbourg, P. (ed.) (1999): Collaborative Learning. Cognitive and Computational Approaches. Elsevier Science Ltd. 1-19.
2. Bentley, R., Appelt, W., Busbach. U., Hinrichs, E., Kerr, D., Sikkel, S., Trevor, J. and Woetzel, G. (1997) Basic Support for Cooperative Work on the World Wide Web. Int. J. of Human-Computer Studies 46(6) 827-846.
3. Caballé S., Xhafa, F., Daradoumis, T. and Marquès, J.M. (2004) Towards a Generic Platform for Developing CSCL Applications Using Grid Infrastructure. In: Proc. of the CLAG/CCGRID'04, Chicago, USA.
4. Rouillard, John P. (2004) Real-time log file analysis using the Simple Event Correlator (SEC). LISA 2004 Conference
5. Xhafa, F., Caballé, S., Daradoumis, Th. and Zhou, N. (2004). A Grid-Based Approach for Processing Group Activity Log Files. In: proc. of the GADA'04, Cyprus.
6. Foster, I. and Kesselman, C. (Eds) (1999) The Grid 2e, 2nd Edition Blueprint for a New Computing Infrastructure. Morgan-Kaufman
7. Open University of Catalonia http://www.uoc.edu (web page as of October 2005).
8. Felix Salfner, Steffen Tschirpke, Miroslaw Malek, Humboldt-University Berlin (2004) Comprehensive Logfiles for Autonomic Systems. 18th International Parallel and Distributed Processing Symposium (IPDPS'04) - Workshop 11
9. Genady Grabarnik, Abdi Salahshour, Balan Subramanian, Sheng Ma, IBM T.J. Watson Research Center (2004) *Generic Adapter Logging Toolkit*. International Conference on Autonomic Computing (ICAC'04)
10. Autonomic Computing Toolkit : http://www-106.ibm.com/developerworks/autonomic/overview.html (web page as of February 2005)
11. Hyades: http://eclipse.org/hyades/ (web page as of February 2005)
12. J. E. Hopcroft, R. Motwani, and J. D. Ullman, Introduction to automata theory, languages, and computation, 2nd ed., Addison-Wesley, 2001.
13. R.A. Baeza-Yates and G.H. Gonnet (1996). Fast Text Searching for Regular Expressions or Automaton Searching on Tries. Journal of the ACM, 43(6):915-936.
14. Goux, J.P., Kulkarni, S., Linderoth, J. and Yoder, M. (2000): An enabling framework for master-worker applications on the computational Grid. In 9th IEEE International Symposium on High Performance Distributed Computing (HPDC'00). IEEE Computer Society.
15. Globus: http://www.globus.org (web page as of December 2004).
16. Planetlab: http://www.planet-lab.org (web page as of February 2004).

C. Paniagua, F. Xhafa, S. Caballé and T. Daradoumis
Open University of Catalonia, Department of Information Sciences

Research Project Focus: GridCole

Gridcole is a new tailorable collaborative learning system that provides guidance to the participants of collaborative learning scenarios defined by educators and enables the use of tools requiring supercomputing capabilities or specific hardware resources for the support of such scenarios. This article provides an overview of the system and describes a sample scenario that can be supported by the system in order to better illustrate its capabilities.

1. Introduction

Computer Supported Collaborative Learning (CSCL) [17] is a mature research field of increasing interest in recent years. As a result of the efforts made in this research domain, many systems have been developed in order to promote and support collaborative methods of learning. These systems, which are typically called collaborative learning systems, are software applications that usually include a set of tools in order to support a given collaborative learning scenario.

One of the most desired features in learning systems in general is tailorability. A computer system is said tailorable if it provides users with some means to modify its functionalities to better suit their needs [20]. In this sense, tailorable collaborative learning systems typically enable easy integration of suitable tools in order to support the realization of new learning scenarios. Examples of this kind of systems include DARE [4], Symba [1] and Cure [9]. Unfortunately, the tailorable collaborative learning systems that can be found in the literature have two important drawbacks.

On the one hand, current systems do not enable the integration of tools requiring supercomputing capabilities or specific hardware devices. However, this type of tools is essential in order to support collaborative learning scenarios in many areas such as engineering, medicine and natural sciences. This idea is supported by the fact that there are a number of non-tailorable collaborative learning systems in which this kind of tools is employed. For example, Covase [15] employs supercomputing capabilities in order to generate a virtual reality environment in which students can interact with 3D complex models related to fluid dynamics. Instead, Pearl [18] accesses to specific hardware devices such as signal pattern generators and oscilloscopes so that students can collaboratory work with a remote laboratory of electronics.

On the other hand, these systems cannot interpret collaboration scripts. A collaboration script is a set of instructions prescribing, among other things, the sequence of activities to be performed by students, as well as the way they should collaborate [6] in order to achieve predefined pedagogical objectives. Following [6], the realization of learning situations according to well-defined collaboration scripts may increase the effectiveness of learning. This can be achieved in learning systems by interpreting collaboration scripts in order to guide students through the sequence of activities to be performed in a given scenario [16] while providing suitable tools and documents to support each of them. Universanté [6] is a good example of a non-tailorable collaborative learning system that provides guidance to students.

Within this framework, we have proposed Gridcole [2,3], a new tailorable collaborative learning system that overcomes these drawbacks. Gridcole exploits grid service technology [7] in order to enable the integration of tools that are not limited in terms of supercomputing and specific hardware needs. Furthermore, it is able to interpret collaboration scripts described by the educator using the IMS-LD (IMS Learning Design) [13] specification.

2. Gridcole approach

Gridcole combines grid service technologies and IMS-LD specification in order to overcome the drawbacks of current tailorable collaborative learning systems.

The service-oriented computational grid is employed in Gridcole as an infrastructure that allows multiple organizations to share a large pool of software tools that may be integrated in order to support all sorts of either individual or collaborative learning scenarios. Such tools are shared as presentation-oriented grid services. More specifically, the business logic of the tool is offered as a service that is compliant to the OGSI (Open Grid Services Infrastructure) [8] specification while the presentation logic is offered as Java-based clients that can be automatically downloaded according to the JNLP (Java Network Launching Protocol) [22] specification. This presentation-oriented grid services sharing approach, which is similar to that employed by other grid-based systems such as Telescience [21], enables the use of tools that are not limited in terms of supercomputing or specific hardware requirements within the context of Gridcole.

In addition, the IMS-LD specification is employed along with the extension presented in [10] in order to specify the collaboration scripts that are to be supported by the system. Furthermore, the use of IMS-LD is combined with the IMS-LRM (IMS Learning Resource Metadata) specification [12] in order to enable the description of the tools that should be employed to support a given scenario, and with the IMS-CP (IMS Content Packaging) [14] specification in order to include all the information required by the system to integrate the actual tools that will be employed in a scenario. IMS-LD documents are thus employed in Gridcole in two different ways. First, as collaboration scripts that can be interpreted by the system in order provide guidance to the participants of the specified scenario. Second, as tailoring scripts describing the tools that must be integrated by the system in order to properly support each activity defined in the collaboration script.

3. System overview

Educators can describe a collaborative learning scenario, listing the activities to be performed, the roles of the users, and the documents and tools that can be used. This information is formalized in an IMS-LD document packaged as a unit of learning according the IMS-CP specification.

Units of learning can be complete or incomplete. A complete unit of learning includes all the information required by the system in order to integrate the actual tools to be used during the realization of the scenario. On the contrary, an incomplete unit of learning does not include such information, but a generic description of each of the tools required for the scenario according to the IMS-LRM specification. Besides, both types of units of learning may include the documents to be employed in the situation or, alternatively, the location of such documents.

Gridcole can only realize scenarios from complete units of learning. However, the system can be employed to generate complete units from incomplete ones, by searching the tools offered in the grid that best match the descriptions included in the incomplete unit of learning. The list of tools found by the system is then provided to the educator so that he can choose those tools that he considers more suitable for the scenario to be realized. Next, the system generates a complete unit including the information that enables the integration of the tools selected by the educator. The complete unit of learning of learning is then stored in the system for later use.

When the educator decides to start the realization of a scenario described in a complete unit of learning, the system provides him with the list of users registered in the system as well as of the different roles that are defined for participants in the collaboration script. The educator can then specify the users that are allowed to join the realization as well as the roles that each of them will play. Once this is done, selected users can join the system in order to join the realization of the situation.

During the realization of a situation, Gridcole interprets the corresponding IMS-LD document in order to inform each participant about the activity that he should be performing in every moment according to the collaboration script as well as about the tools and documents that have been defined for the support of the activity. If the participant decides to work with one of the documents, the system will automatically download it for him. Instead, if the participant decides to use a tool, Gridcole automatically installs and launches its client in the user's machine so that he can start using the tool immediately.

4. Architecture outline

Gridcole architecture, as presented in Figure 1, consists of a web portal, the Gridcole client, and a Learning Flow Engine (LFE).

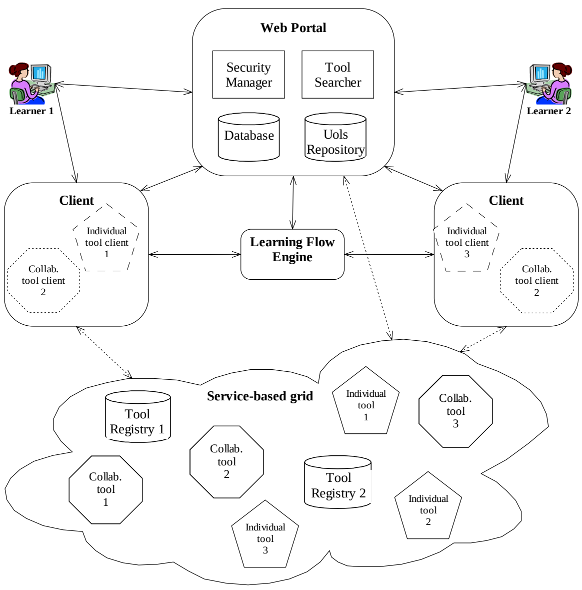

Figure 1. Outline of Gridcole architecture.

The web portal authenticates users and provides them with a single access point to all system functionalities. Within the web portal, a *database* keeps administrative information about users, but also information about the state of the units of learning being executed. The *unit repository* stores all units of learning, both incomplete and complete, that are available in the system. The *tool searcher* seeks for available tools in the grid, according to the needs stated by educators in an incomplete unit of learning. To do so, it queries *tool registries*, where service providers advertise their services including descriptions according to the IMS-LRM specification. Finally, a *security manager* provides users with suitable credentials in order to access tools offered as secure grid services.

The LFE interprets the IMS-LD file corresponding to complete units of learning, and determines the activities to be carried out at each moment and the users participating in them. It also determines which tools and documents can be used in each activity, and if necessary, connects tool clients to the appropriate tool service instances.

The Gridcole client provides a graphical interface, similar to a desktop, and is instructed by the LFE on which activity is being performed by the user, and which tools and documents are available. If the user selects a document, the Gridcole client downloads it from the appropriate location in the grid. If the user selects a tool, the Gridcole client transparently downloads and installs the tool client, which is then connected to the tool service part, being executed in the grid.

5. Prototype

The Gridcole system is currently under development. However, a prototype has already been built in order to show the feasibility of the proposed system. The prototype is based on the learning workflow engine and the engine client developed by the Coppercore project [5]. Both elements have been conveniently modified in order to enable users to employ grid-based tools during the realization of collaborative learning scenarios.

This prototype can only be employed in order to realize collaborative learning scenarios described in complete units of learning. Furthermore, the user is responsible for carrying out some operations that should be performed by the system in order to start the realization of the scenario. However, once this is done, the prototype can be used in order to participate in the realization of scenarios without any limitation.

6. Sample collaborative learning scenario

In order to further illustrate the use of Gridcole, this section shows how our system can be employed to support a real collaborative learning scenario requiring tools with special hardware needs. This sample scenario has been worked out for a course on Computer Architecture, at the University of Valladolid.

In this scenario, and according to the educational motivation described in [19], students play the role of consultants that have to make a recommendation on a computing system to a fictitious (but realistic) customer, played by the teacher. In the first part of the course, the benchmarking topic is studied. In the traditional realization of the course, students benchmark the only three existing architectures at their laboratory, and have to recommend one of them. With the use of Gridcole, any computer architecture existing in the grid (of course not necessarily restricted to the University of Valladolid) can be benchmarked. Furthermore, with Gridcole students can be guided through a sequence of activities planned by the educator, improving efficiency and collaboration in their laboratory assignments.

The collaboration script corresponding to this scenario states the following activities:

1. Individually study customer requirements and propose a workload model.
2. Individually study the documentation on available machines and benchmarks, and propose an evaluation plan.
3. Individually benchmark available machines and propose a recommendation.
4. In pairs, determined by the system from the IMS-LD script, discuss the workload model and propose an agreed model.
5. In pairs, discuss the evaluation plan, and propose an agreed plan.
6. In pairs, if necessary repeat experiments, or recombine results, and propose an agreed machine recommendation.

The realization of this scenario requires some documents (customer description, machine features, benchmark descriptions…) and some tools (benchmarking, document editing, chatting…) to be available to the users. Besides, the system realizing this scenario must be able to allow accessing to specific hardware resources (the machines to benchmark). Moreover, it must be able to interpret and execute collaboration scripts.

In order to show that Gridcole system can be actually employed to realize this scenario, we have developed the benchmarking tool and the chat tool required, according to the specifications mentioned in section 2. Furthermore, we have formalized the collaboration script outlined above using the IMS-LD specification and packaged it in a complete unit of learning. These elements have been employed along with the Gridcole prototype in order to support the realization of the scenario in two experiments with real users. More specifically, 4 former students of the course on Computer Architecture participated in the first experiment while 8 actual students of the course participated in the second. According to the feedback provided by participants in questionnaires, 75% of them considered that the system was helpful or very helpful to carry out the scenario in collaboration with their partners. Figure 2 shows a snapshot of Gridcole prototype during the realization of the benchmarking activity included in the scenario.

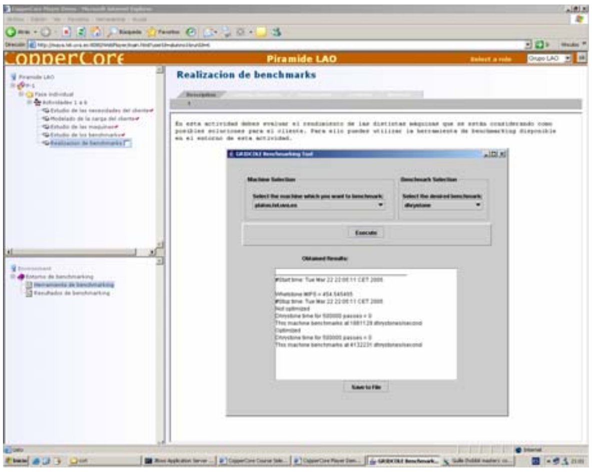

Figure 2. Snapshot of Gridcole client along with the benchmarking tool as seen by students during the realization of the benchmarking activity.

7. Conclusions

This article provided an overview of Gridcole, a system that can be easily tailored by educators in order to support the realization of collaborative learning scenarios designed by themselves. To do so, educators can provide a script specifying the sequence of learning activities to be performed by students as well as the documents and tools required to support them. Gridcole can then search for these tools in a service-based grid in order to make them available to users whenever it is required. Significantly, these tools are not limited in terms of access to supercomputing capabilities or specific hardware resources. Furthermore, Gridcole can guide students during the realization of the collaborative learning scenario according to the sequence of activities defined in the script.

Upcoming work includes the completion of Gridcole prototype, the enhancement of Gridcole's tool search capabilities using ontologies [23], and the integration of the system with the Collage [11] IMS-LD editor of collaborative learning scenarios. In the long term we plan to study the possibility of integrating tools based on P2P technology in Gridcole. Furthermore, scheduling facilities, which are usually employed in grid computing, will be studied in order to provide collaborative tools according to predefined Quality of Service (QoS) parameters.

References

[1] M. L. Betbeder, P. Tchounikine. "Symba, a tailorable framework to support collective activities in a learning context". Proceedings of the 9th International Workshop on Groupware (CRIWG 2003). Grenoble, France, 2003.

[2] M. L. Bote-Lorenzo, D. Hernández-Leo, Y. A. Dimitriadis, J. I. Asensio-Pérez, E. Gómez-Sánchez, G. Vega-Gorgojo, L. M. Vaquero-González. "Towards reusability and tailorability in collaborative learning systems using IMS-LD and grid services". International Journal on Advanced Technology for Learning, pp. 129-138, vol. 1 (3) (2004).

[3] M. L. Bote-Lorenzo, L. M. Vaquero-González, G. Vega-Gorgojo, Y. A. Dimitriadis, J. I. Asensio-Pérez, E. Gómez-Sánchez, D. Hernández-Leo. "A tailorable collaborative learning system that combines OGSA grid services and IMS-LD scripting". Proceedings of the 10th International Workshop on Groupware (CRIWG 2004). San Carlos, Costa Rica, 2004.

[4] G. Bourguin, A. Derycke. "Integrating the CSCL activities into virtual campuses: foundations of a new infrastructure for distributed collective activities". Proceedings of the 1st European Conference on Computer Supported Collaborative Learning (EuroCSCL 2001). Maastricht, The Netherlands, 2001.

[5] CopperCore Project: http://coppercore.org/.

[6] P. Dillenbourg. "Over-scripting CSCL: the risks of blending collaborative learning with instructional design". In: Three worlds of CSCL. Can we support CSCL, ed. P. A. Kirschner. Heerlen, Open Universiteit Nederland, pp. 61-91 (2002).

[7] I. Foster, C. Kesselman. The Grid 2: blueprint for a future computing infrastructure, San Francisco, CA, USA: Morgan Kaufmann Publishers (2004).

[8] Global Grid Forum (GGF). "Open Grid Services Infrastructure (OGSI) version 1.0". Technical Specification. Jun 2003.

[9] J. M. Haake, T. Schümmer, A. Haake, M. Bourimi, B. Landgraf. "Two-level tailoring support for CSCL". Proceedings of the 9th International Workshop on Groupware (CRIWG 2003). Grenoble, France, 2003.

[10] D. Hernández-Leo, J. I. Asensio-Pérez, Y. Dimitriadis. "IMS Learning Design Support for the Formalization of Collaborative Learning Patterns". Proceedings of the 4th International Conference on Advanced Learn-

ing Technologies (ICALT 2004). Joensuu, Finland, 2004.

[11] D. Hernández-Leo, E. D. Villasclaras-Fernández, I. M. Jorrín-Abellán, J. I. Asensio-Pérez, Y. Dimitriadis, I. Ruiz-Requies, B. Rubia-Avi. "COLLAGE, a collaborative learning design editor based on patterns (accepted for publication)". Journal of Educational Technology & Society, Special Issue on Learning Design (2006).

[12] IMS Global Learning Consortium. "IMS Learning Resource Metadata information model v1.2.1". Technical Specification. Sep 2001.

[13] IMS Global Learning Consortium. "IMS Learning Design Information Model v1.0". Technical Specification. Jan 2003.

[14] IMS Global Learning Consortium. "IMS Content Packaging Information Model v1.1.3". Technical Specification. Jun 2003.

[15] N. Jensen, S. Seipel, W. Nejdl, S. Olbrich. "COVASE: Collaborative visualization for constructivist learning". Proceedings of the International Conference on Computer Support for Collaborative Learning (CSCL 2003) . Bergen, Norway, 2003 .

[16] P. Jermann, A. Soller, A. Lesgold. "Computer software support for collaborative learning". In: What we know about CSCL in higher education, eds. J. W. Strijbos, P. Kirschner, R. Martens. Amsterdam, The Netherlands: Kluwer, pp. 141-146 (2004).

[17] T. Koschmann. CSCL: theory and practice of an emerging paradigm, Malwah, NJ, USA: Lawrence Erlbaum (1996).

[18] J. M. Martins Ferreira, G. R. C. Alves, R. Costa, N. Hine. "Collaborative learning in a web-accesible workbench". Proceedings of the 8th International Workshop on Groupware (CRIWG 2002). La Serena, Chile, 2002.

[19] A. Martínez-Monés, E. Gómez-Sánchez, Y. Dimitriadis, I. M. Jorrín-Abellán, B. Rubia-Avi, G. Vega-Gorgojo. "Multiple case studies to enhance project-based learning in a computer architecture course". IEEE Transactions on Education, pp. 472-481, vol. 48 (3) (2005).

[20] A. Morch. "Three levels of end-user tailoring: customization, integration and extension". Proceedings of the 3rd Decennial Aarhus Conference. Aarhus, Denmark, 1995.

[21] S. Peltier, M. E. Martone, S. Lamont, A. Lin, D. Lee, T. Molina, L. Dai, M. Wong, S. Mock, M. H. Ellisman. "The Telescience portal for advanced tomography applications". Journal of parallel and distributed computing, pp. 539-550, vol. 63 (5) (2003).

[22] Sun Microsystems. "Java Network Launching Protocol & API specification version 1.5". Technical Specification. May 2001.

[23] G. Vega-Gorgojo, M. L. Bote-Lorenzo, E. Gómez-Sánchez, Y. A. Dimitriadis, J. I. Asensio-Pérez. "Semantic search of learning services in a grid-based collaborative system". Proceedings of the 2nd International Workshop on Collaborative and Learning Applications of Grid Technology and Grid Education (CLAG+Grid.Edu 2005), held in conjunction with the International Symposium on Cluster Computing and the Grid 2005 (CCGRID 2005). Cardiff, UK, 2005.

M.L. Bote-Lorenzo, E. Gómez-Sánchez, G. Vega-Gorgojo, Y. Dimitriadis, and J.I. Asensio-Pérez
GSIC/EMIC Research group
(http://gsic.tel.uva.es)

University of Valladolid, Spain

Technology Watch

This section presents Technologies, Specifications and Standards related to the e-Learning and GRID world. A brief description will be given together with a set of references to "must read" articles and documents.

The Globus® Toolkit has become the *de facto* GRID middleware standard. In this section we will present the Globus Project and the Globus Toolkit.

Globus Alliance

The *Globus Alliance* is a group of organisations and individuals that work behind the scenes of the 'Grid'. Their work is mainly focused on the developing of fundamental technologies to support the applications based on Grid.

The *Globus Consortium* is a non-profit organisation formed by global computing leaders who support the Globus Toolkit, the de facto standard for open source grid computing infrastructure. With the full support of leading enterprise hardware and software vendors, the original pioneers of Grid, and the open source Grid development community - The Globus Consortium will leverage its broad base of participants to further accelerate the evolution of Grid in the enterprise.

This group's main activity is to work on research and development in the standardisation field of the Grid infrastructure and applications.

The activity began in 1995 and the group is based are Argonne National Laboratory [1], the University of Southern California's Information Sciences Institute, the University of Chicago [3], the University of Edinburgh [4], the Swedish Center for Parallel Computers [5], and the National Center for Supercomputing Applications (NCSA) [6].

The group produces open-source software that is central to the building of applications based on Grid. It offers software and toolkits that allow the creation of new applications in an easy way for developers.

Apart from the core group, the Globus Alliance has a programme to get new contributions from relevant actors. The Academic Affiliates are spread all over the world and counts on research and educational institutions from Asia, Europe and America.

About the Globus® Toolkit

The open source Globus Toolkit is fundamental in enabling technologies for the 'Grid,' letting people share computing power, databases, and other tools securely online across corporate, institutional, and geographic boundaries.

The toolkit includes software services and libraries for resource monitoring, discovery, and management, plus security and file management. In addition to being a central part of science and engineering projects that sum nearly a half-billion dollars internationally, the Globus Toolkit is a substrate on which leading IT companies are building significant commercial Grid products.

The toolkit includes software for: security, information infrastructure, resource management, data management, communication, fault detection, and portability. It is packaged as a set of components that can be used either independently or together to develop applications. Every organization has unique modes of operation, and collaboration between multiple organizations is hindered by incompatibility of resources such as data archives, computers, and networks. The Globus Toolkit was conceived to remove obstacles facilitating the collaboration among different organisations. Its core services, interfaces and protocols allow users to access remote resources as if they were located within their own machine room while simultaneously preserving local control over who can use resources and when.

The Globus® Toolkit has grown through an open-source strategy similar to the Linux operating systems, and distinct from proprietary attempts at resource-sharing software. This encourages broader, more rapid adoption and leads to greater technical innovation, as the open-source community provides continual enhancements to the product..." [adapted from Reference 7]

The Globus Toolkit is "an open source software toolkit used for building enterprise-level Grid systems and applications. Freely available in open source format on the Web, the Globus Toolkit provides applications for security, information infrastructure, resource management, data management, communication, fault detection, portability, and more.

Since 1990s several researchers groups have contributed to Globus Toolkit code. Currently, the last version that has been produced is GT4. This version includes all the necessary components and facilities to enable the creation of an enterprise Grid.

The key features that the GT4 presents may be summarised as follows:

- GT4 complies with the latest Web Services Interoperability Organization (WS-I) web services standards, which provides maximum interoperability between different environments.

- GT4 includes initial support for important authorization standards, including Security Markup Language (SAML) and Extensible Access Control Markup Language (XACML). These provide business with a foundation for building secure web-services enabled Grid infrastructure.
- GT4 implements the Web Services Resource Framework (WS-RF) and Web Services Notification (WS-N) specifications, which are emerging standards in OASIS backed by major vendors for web services enablement of Grid and resource management systems.
- GT4 features sophisticated authorisation and security capabilities. Globus has always been diligent in Grid security, and GT4 is also "enterprise ready" from a security perspective.

The Globus® Toolkit is available from http://www.globus.org/toolkit/. The distribution of the code is done under an open source license that gives a certain freedom to use and distribute it.

The Globus Toolkit evolves towards the same direction as the Open Grid Services Architecture (OGSA) offering solutions that follow the OGSA principles to create new services. The Open Grid Services Architecture (OGSA) represents an evolution towards a Grid system architecture based on Web services concepts and technologies. OGSA is a product of the Grid community at large, and it has a major focal point in the Global Grid Forum (GGF). Members of the Globus Alliance have made significant contributions to the development of OGSA.

Learn more about Globus

The description presented in this article is just a small piece of content of the Globus world. The Globus alliance has produced very different documents to let developers know about the tools (not only what it is available but also how to use it, how to contribute, etc.).

In order to learn more about this group, their activities and their Tool Kit releases it is necessary to visit www.globus.org. In order to get started and now more about the idea the following articles are a good reference to look at:

- **Globus Toolkit Version 4: Software for Service-Oriented Systems**. I. Foster. *IFIP International Conference on Network and Parallel Computing*, Springer-Verlag LNCS 3779, pp2-13,2005.
- **The Anatomy of the Grid: Enabling Scalable Virtual Organizations**. I. Foster, C. Kesselman, S. Tuecke. *International J. Supercomputer Applications*, 15(3), 2001.
- **The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration**. I. Foster, C. Kesselman, J. Nick, S. Tuecke, Open Grid Service Infrastructure WG, Global Grid Forum, June 22, 2002.
- **State and Events for Web Services: A Comparison of Five WS-Resource Framework and WS-Notification Implementations**. M. Humphrey, G. Wasson, K. Jackson, J. Boverhof, M. Rodriguez, Joe Bester, J. Gawor, S. Lang, I. Foster, S. Meder, S. Pickles, and M. McKeown, 4th IEEE International Symposium on High Performance Distributed Computing (HPDC-14), Research Triangle Park, NC, 24-27 July 2005.
- **From Open Grid Services Infrastructure to WS-Resource Framework: Refactoring & Evolution**. K. Czajkowski, D. Ferguson, I. Foster, J. Frey, S. Graham, T. Maguire, D. Snelling, S. Tuecke, March 5, 2004.
- **Open Grid Services Infrastructure (OGSI) Version 1.0**. S. Tuecke, K. Czajkowski, I. Foster, J. Frey, S. Graham, C. Kesselman, T. Maguire, T. Sandholm, P. Vanderbilt, D. Snelling; Global Grid Forum Draft Recommendation,6/27/2003.
- **The Globus Project: A Status Report.** I. Foster, C. Kesselman. *Proc. IPPS/SPDP '98 Heterogeneous Computing Workshop*, pp. 4-18,1998.

These can be found among other interesting references and articles on the web site of the Globus Project.

References

[1] http://www.anl.gov/

[2] http://www.isi.edu/

[3] http://www.uchicago.edu/

[4] http://www.ed.ac.uk/

[5] http://www.pdc.kth.se/

[6] http://www.ncsa.uiuc.edu/

[7] http://www.globus.org

Blanca Jordan
ATOS Origin

News

Semantic Grid: a Call for Papers for a Special Issue of FGCS.

5 October 2005

The Internet and WWW are milestones of information sharing. IT professionals are in the critical historical stage of creating the new brilliance a new interconnection environment. Semantics and knowledge have become the kernel issues in developing the new environment.

To face these issues, Future Generation Computer Systems (The International Journal Of Grid Computing: Theory, Methods And Applications) will release a Special Issue on Semantic Grid and Knowledge Grid. This special issue focuses on systems and applications on following topics:

- Semantics and Semantic Grid (Semantic Computting; Semantic Interoperability; Semantic Capture; Automatic Semantic Annotation; Semantic Web Services; Theory, Model and Applications of Semantic Grid)
- Knowledge and Knowledge Grid (Theory, Model and Applications of Knowledge Grid; Distributed Knowledge Management; Information, Knowledge and Service Integration; Ontology; Knowledge Flow; Knowledge Discovery; Internet-based Knowledge Engineering and Software Engineering.
- Advanced Computing Model (Autonomous Computing and Cooperative Networking; Internet-based Intelligence and Applications; P2P and Grid Computing; Self-organized Intelligence and Service-oriented Computing
- e-Science, e-Culture, e-Business, e-Learning, e-Government

The deadline for Submission is November 30, 2005. For more information about this special issue, see: http://www.elsevier.com/locate/future

Grid – the key to scientific collaboration

22 September 2005

A joint UNESCO-ROSTE and CERN event on ***grid*** computing, 'Grid - the Key to Scientific Collaboration', took place in Geneva, Switzerland, from 28 to 29 September.

Grid computing, designed as an open, distributed computing paradigm, has rapidly moved to the centre of attention as a key future-oriented technology. Based on the collaboration-fostering and research-enabling role of the ***grid***, CERN and UNESCO-ROSTE are taking the opportunity to invite current and future ***grid*** participants, universities and research institutions to a ***grid*** event.

Through presentations by key ***grid*** protagonists, including some from CERN who are involved in the European Commission funded Enabling Grids for E-sciencE (EGEE) project, participants were able to learn about the capabilities of the ***grid***, opportunities to leverage their research work, and how to participate in international projects.

For more information:
http://www.hp.com/events/*grid*

Global Grid Forum 14: Grids for Education and Outreach

28 June 2005

A session devoted to the education community was held in GGF14. The session consisted of four talks addressing using Grids to enhance education.

Randy Ruchti discusses **QuarkNet,** which is a collaboration between Fermilab and the particle physics Grid activities (open Science Grid) to use the Grid to bring the excitement of physics research to schools. This example of Grids in support of STEM (Science, Technology, Engineering and Mathematics) is built on by Bill Frascella who was director of ESIE (Elementary Secondary and Informal Education) component of EHR at NSF. Frascella has returned to Indiana University where he founded the **Indiana Mathematics Initiative for middle schools**. He explained how Grids can help bring resources such as virtual laboratories and science expertise to schools across the nation.

The third talk came from Al Kuslikis of **AIHEC** (American Indian Higher Education Consortium) who described how Grids can help integrate Tribal Homelands into a broad community and in particular how e-Science is important for the Tribal Colleges. The final talk by James Turner of Virginia Tech described the importance of Grids in developing technology, technologists and scientists in Africa. The initial effort was built around the netork of institutes forming the **African Institute for Mathematical Sciences** (AIMS)

Presentations are availble at:
http://www.ggf.org/GGF14/GGF14Education.html

When	What	Where
25 - 28 October 2005	**mLearn 2005 - 4th World conference on mLearning Conference theme: Mobile technology: The future of learning in your hands** This annual conference is the key research and networking event for researchers, strategists, educators, technologists and practitioners from all over the world. Previous mLearn conferences have attracted participants from more than 60 countries, and is, therefore, the world's largest conference on mLearning and emerging ambient technologies. For more information: http://www.mlearn.org.za/	**Cape Town, South Africa**
10 –11 November 2005	**ECEL 2005: The 4th European Conference on e-Learning** The European Conference on e-Learning (ECEL) welcomes research in all aspects of this growing area of education. Actively supported by both academics and practitioners this conference provider an opportunity for networking and sharing advancing theories and practice. For more information: http://www.academic-conferences.org/ecel/ecel-home.htm	**Amsterdam, Netherlands**
Nov 30- Dec 2 2005	**11th International Conference on Technology Supported Learning and training, ONLINE EDUCA BERLIN 2005** Meeting the networking needs of the international e-learning and distance education industry, the annual Online Educa Berlin conference is the **key networking venue for strategists and practitioners from all over the world**. The **Kaleidoscope Network of Excellence** in technology-enhanced learning and **eLIG** (European eLearning Industry Group) will bring together industry and academia to discuss the creation of a good practice model for linking scientific research and innovation. Also put under the microscope will be industrial exploitation in technology-enhanced learning for shared understanding and consensus on the priorities of basic research in the field. http://www.online-educa.com/en/	**Berlin, Germany**
November 28-30, 2005	**IEEE International Workshop on Wireless and Mobile Technologies in Education WMTE 2005** In today's globally competitive environment, effective wireless and mobile technologies allows new opportunities for users and learners to be intensely connected. Advanced learning theories advocate many social aspects of learning such as collaborative learning, communities of practice, internalization of social process, participation in joint activity, and situated learning. WMTE 2005 will bring together researchers, academics and industry practitioners who are involved or interested in the design and development of Wireless and Mobile Learning Technologies. Understanding of the challenges faced in providing technology tools to support the learning process and ease the creation of instruction material using mobile technologies will help building a direction for further research and implementation work in the ubiquitous learning society. For more information: http://lttf.ieee.org/wmte2005/netscape/index.html	**The University of Tokushima, Japan**

When	What	Where
November 28-29, 2005	**1st International Conference on Semantics, Knowledge and Grid** The Internet and World Wide Web are milestones of information sharing. IT professionals are in the critical historical stage of creating the new brilliance --- a new interconnection environment. Semantics and knowledge have become the kernel issues in developing the new environment. The 1st International Conference on Semantics, Knowledge and Grid (SKG 2005) is to bring together researchers and practitioners in areas of knowledge and intelligence, semantics, and grid computing to share their visions, research achievements and solutions to real problems, to attack the challenge issues, and to establish worldwide cooperative research and development. For more information: http://kg.ict.ac.cn/SKG2005/	**Beijing, China**
December, 5 – 8, 2005	**e-Science 2005 – International Conference on e-Science and Grid technologies** The next generation of scientific research and experiments will be carried out by communities of researchers from organizations that span national boundaries. These activities will involve geographically distributed and heterogeneous resources such as computational systems, scientific instruments, databases, sensors, software components, networks, and people. Such large-scale and enhanced scientific endeavors, popularly termed as e-Science, are carried out via collaborations on a global scale. Grid computing has emerged as one of the key computing paradigms that enable the creation and management of Internet-based utility computing infrastructure, called Cyberinfrastructure, for realization of e-Science and e-Business at the global level. Several national and international projects around the world have been initiated to carry out research and innovation activities that transform the vision of e-Science and Grid computing into reality. The e-Science 2005 conference, sponsored by the IEEE Computer Society Technical Committee for Scalable Computing (TCSC) and Society for Industrial and Applied Mathematics, is designed to bring together leading international and interdisciplinary research communities, developers, and users of e-Science applications and enabling IT technologies. The conference serves as a forum to present the results of the latest research and product/tool developments, and highlight related activities from around the world. For more information: http://www.gridbus.org/escience/	**Melbourne , Australia**
December 14-16, 2005	**IADIS International Conference. Cognition and Exploratory Learning in Digital Age (CELDA 2005)** The IADIS CELDA 2005 conference aims to address the main issues concerning with the evolving learning processes and supporting pedagogies and applications in digital age. There have been huge advancements in both cognitive psychology and computing that have affected the educational arena. The convergence of these two disciplines has never been faster before and this marriage has affected the very basis of the academia. Paradigms such as just-in-time learning, constructivist approaches, student-centered learning and collaborative approaches have emerged, and are being supported by technological advancements such as simulations, virtual reality and multi-agents systems to name a few. This merger has created both opportunities and areas of serious concerns. This conference aims to cover both technological as well as pedagogical issues related to these developments. Main tracks have been identified. However innovative contributions that do not fit into these areas will also be considered as long as they are directly related to the overall theme of the conference. For more information: http://www.iadis.org/celda2005/	**Porto, Portugal**

When	What	Where
February 26-28, 2006	**IADIS International Conference. Web Based Communities 2006** The mission of this conference is to publish and integrate scientific results and act catalytically to the fast developing culture of web communities. The conference invites original papers, review papers, technical reports and case studies on WWW in particular the emerging role of so-called WWW-Based Communities. For more information: http://www.iadis.org/wbc2006/index.asp	**San Sebastián, Spain**
May 16-19, 2006	**Collaborative and Learning Applications of Grid Technology CLAG 2006** CLAG 2006 is the third edition of a successful international workshop devoted to research on the application of grid technologies within the context of e-learning and collaboration. This workshop will be held in conjunction with CCGrid 2006, which is one of the most important international conferences in grid computing. CLAG 2006 provides researchers with a good opportunity to present their latest works and to share experiences dealing with the topics covered by the workshop. Extended versions of the best papers presented in CLAG 2006 will be peer-reviewed for publication in a special issue of the international journal of Future Generation of Computing Systems. For more information: http://gsic.tel.uva.es/clag/clag2006.html/	**Singapore, Malaysia**