Proc. 3st Int. Conf. "Information Technologies and Management", April 14-15, 2005 (Information Systems Management Institute, Riga, Latvia, 2005) pp. - .

TEACHING CLUSTER COMPUTING AT LATVIAN UNIVERSITY

ALEXEI KUZMIN

Institute of Solid State Physics of Latvian University
Kengaraga Street 8, LV-1010, Riga, Latvia
e-mail: a.kuzmin@cfi.lu.lv

ABSTRACT

The problem of teaching cluster computing is discussed on an example of the course “Introduction to Cluster Computing”, given for the first time at the Latvian University during the first semester of 2004/2005 scholastic year. The course introduces into the modern supercomputing technology - cluster computing and provides with the knowledge on the state-of-the-art and architectures of supercomputers, on the theoretical and practical aspects of supercomputers/clusters applications, parallel programming technologies and their use and realisation on clusters. [*Keywords: teaching cluster computing, high performance computing, parallel computing*]

1. INTRODUCTION

During last years, cluster computing becomes more and more popular and attractive in solving large-scale science, engineering, and commercial problems [1]. Currently clusters support applications ranging from supercomputing and mission-critical software, through web server and e-commerce, to high performance database applications. The “cluster” is a group of interconnected computers working together on a single problem (Figure 1). The ability to join power of separate computing units, represents one of the main advantages of cluster systems and, in fact, differentiates them from just a group of computers working independently. The cluster appears to users as a single computing resource, whose functioning is realized within one of the two paradigms [2]: (i) parallel computer emulation (so-called Beowulf-type systems) and (ii) single system image (SSI). The Beowulf clusters utilise external libraries as MPI [3] or PVM [4] and thus require modification of the application program. In the case of the SSI clusters, the user programs should not be changed, and, the system resources are automatically load-balanced that makes their use more efficient.

Nowadays, cluster computing provides an inexpensive computing resource to educational institutions as universities and high-schools, which don’t need to invest millions of dollars to buy massive parallel computers, known as “supercomputers”, for the purpose of teaching "parallel computing". Today, any one can build a small cluster from commodity-off-the-shelf (COTS) hardware components, obtain free software from the web, and use the cluster in the educational process.

Proc. 3st Int. Conf. "Information Technologies and Management", April 14-15, 2005 (Information Systems Management Institute, Riga, Latvia, 2005) pp. - .

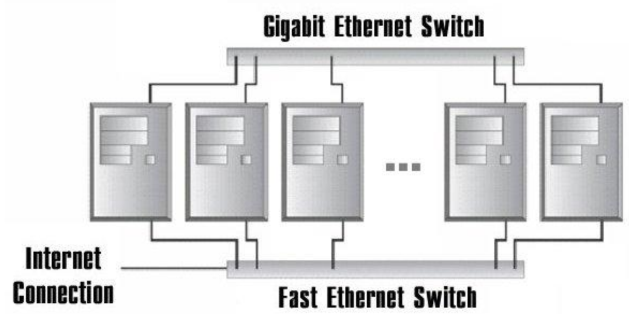

FIGURE 1. STANDARD CLUSTER ARCHITECTURE [1]: A GROUP OF COMPUTERS (NODES) ARE INTERCONNECTED BY TWO LOCAL NETWORKS THROUGH TWO SWITCHES. GIGABIT-TYPE NETWORK IS DEDICATED TO COMMUNICATIONS BETWEEN CLUSTER NODES DURING CALCULATIONS, WHEREAS THE FAST NETWORK SERVES USERS TO CONNECT TO THE CLUSTER.

Many universities and research centers around the world, including those in developing countries, have also used clusters as a platform for high performance computing (HPC). In fact, the cluster systems become firmly a part of the most powerful computers [5]. In November 2004, there have been totally 294 clusters in the Top500 Supercomputers list. Moreover, the fastest cluster systems (BlueGene/L, Columbia and MareNostrum) have occupied the first, second and fourth positions (Figure 2). All this shows general trends in the supercomputing technologies towards cluster computing and indicates todays demands for specialists preparation in the field.

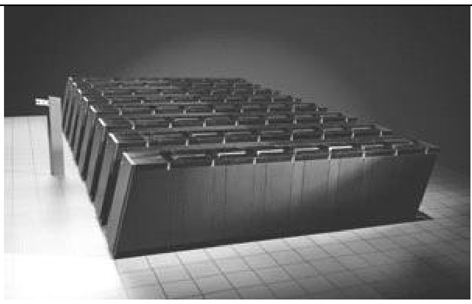

BlueGene/L DD2
IBM 0.7 GHz PowerPC 440
70.72 TFlops, 32768 processors

Columbia
SGI Altix 1.5 GHz
51.87 TFlops, 10160 processors

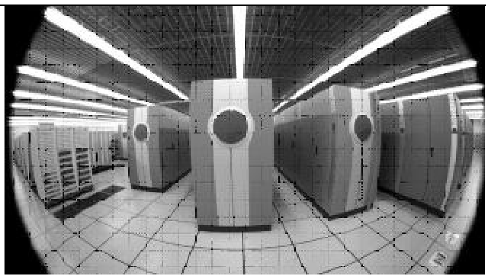

Earth-Simulator
35.86 TFlops, 5120 processors

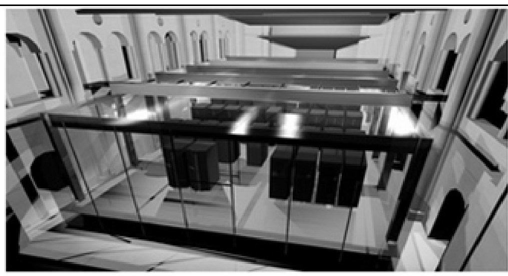

MareNostrum
IBM 0.7 GHz PowerPC 970FX
20.53 TFlops, 3564 processors

FIGURE 2. TOP SUPERCOMPUTER SYSTEMS IN NOVEMBER 2004 [5].

Proc. 3st Int. Conf. "Information Technologies and Management", April 14-15, 2005 (Information Systems Management Institute, Riga, Latvia, 2005) pp. - .

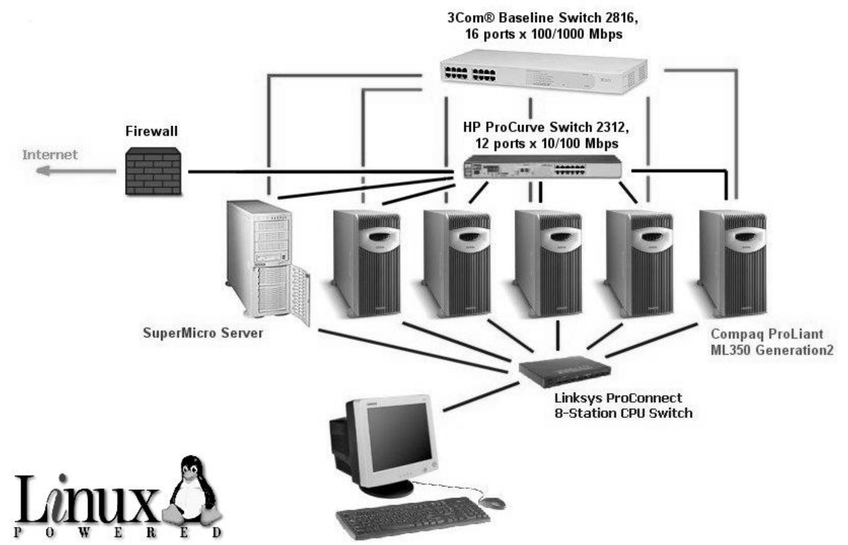

FIGURE 3. LATVIAN SUPERCLUSTER (LASC) ARCHITECTURE IN 2005 [6]. SUPERMICRO SERVER WITH TWO PENTIUM 4-3.06GHZ XEON CPUS PLUS FIVE COMPAQ PROLIANT ML350 G2 SERVERS WITH TWO PENTIUM III-1.3GHZ CPUS ARE INTERCONNECTED BY FAST AND GIGABIT ETHERNET NETWORKS. THEORETICAL PEAK PERFORMANCE IS 19 GFLOPS . TOTAL CLUSTER RESOURCES ARE 25 GB OF OPERATING MEMORY (RAM) AND 1.8 TB OF HARD DISK SPACE.

In 2002 the first HPC Beowulf-type cluster, called Latvian SuperCluster (LASC) [6], has been installed at the Institute of Solid State Physics (ISSP) of the Latvian University (LU). It is used for theoretical calculations in the fields of solid state physics and quantum chemistry by scientists and students, thus being tightly integrated into research and educational processes. The experience of LASC usage has indicated two tendencies. First, the number of LASC users, both researches and students, increases progressively, because of nowadays science demands for the use of more complex theoretical models and for the shortening of research timescale. From another side, most of the users has very weak knowledge of cluster technologies that results in difficulties with the cluster use and inefficient problem solving. Moreover, general trends in the evolution of information technologies and appearance of cluster-type systems on the market and in commercial use require preparation of new specialists with deep knowledge of parallel computing technologies at both hardware and software levels.

In attempt to address this problem, the new lecture course “Introduction to Cluster Computing” [7] has been developed by the author and given for the first time at the Faculty of Education and Psychology of Latvian University within the programme “Computer networks and systems administrator” during the first semester of 2004/2005 scholastic year.

2. DESCRIPTION OF THE LECTURE COURSE “INTRODUCTION TO CLUSTER COMPUTING”

The idea of the course is to introduce students into the field of cluster computing and to provide them with the basic knowledge on the state-of-the-art and architectures of supercomputers, on the theoretical and practical aspects of supercomputers/clusters applications, parallel programming technologies and their use and realisation on clusters.

The course size is 32 hours, among which 24 hours are dedicated to lectures and 8 hours are devoted to practicals and topical reports preparation. The course is supported by the Web-site (Figure 4) [7], where students can obtain information on the course, lecture notes (currently more than 150 pages), a list of recommended literature, a list of internet resources on the subject and useful reference materials. A final exam is intended at the end of the course.

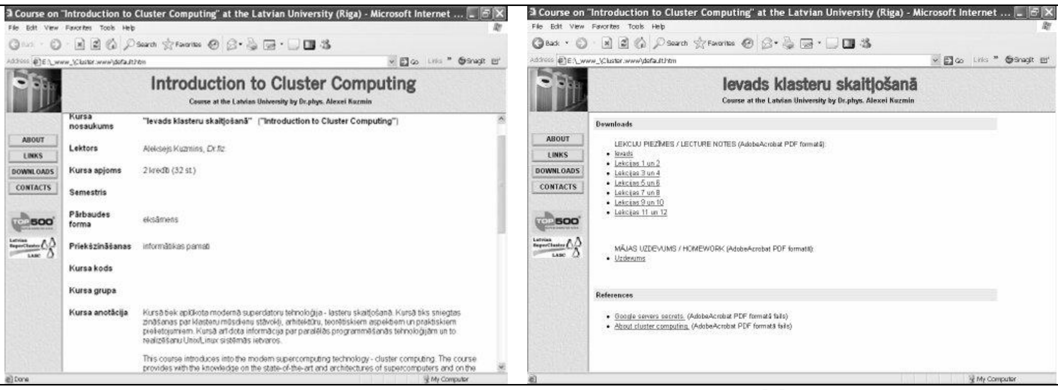

FIGURE 4. EXAMPLE OF WEB-PAGES FROM THE WEB-SITE [7],
DEDICATED TO THE LECTURE COURSE "INTRODUCTION TO CLUSTER COMPUTING".

TABLE 1. CONTENT OF THE COURSE "INTRODUCTION TO CLUSTER COMPUTING" [7].

1. Supercomputers: history, trends, state-of-the-art. Clusters and superclusters.
2. Applications of supercomputers and clusters.
3. Supercomputers architecture: mainframes and clusters.
4. Shared memory systems.
5. Distributed memory systems.
6. Cluster building blocks. Clusters design.
7. System-level middleware.
8. Cluster installation and configuration.
9. Introduction to UNIX/LINUX.
10. Basics of Unix/Linux installation, configuration and administration.
11. Work with Unix/Linux.
12. Practical use of supercomputers.
13. Introduction to parallel programming technologies.
14. Cluster programming using MPI and PVM.
15. Other cluster programming technologies.
16. Parallel algorithms and applications.
17. Monte Carlo modelling on clusters.

The course covers 17 topics, listed in Table 1, which address historical, theoretical, practical and applied aspects of cluster computing. Further, we will describe briefly the topics content. More information on the subject can be found in [1,2,6-9].

The first part of the course (topics 1 and 2 in Table 1) introduces the students into terminology of supercomputing and clustering, gives a historical overview on the supercomputers development, presents their current state-of-the-art and indicates tendencies and trends in the creation of high performance systems based on cluster architecture [1,2,8,9]. Besides, several examples of currently working cluster systems in Latvia [1,6] and abroad [5] are given to the students, and a number of clusters applications is discussed in details.

The course second part (topics 3 to 5 in Table 1) is dedicated to the supercomputers taxonomy [2,8]. The students are introduced into SISD (Single Instruction stream, Single Data stream), MISD (Multiple Instruction stream, Single Data stream), SIMD (Single Instruction stream, Multiple Data stream) and MIMD (Multiple Instruction stream, Multiple Data stream) computer architectures, as well as systems with shared and distributed memories.

The next part (topic 6 in Table 1) is dedicated to the clusters anatomy [10-14]. The students aquire the knowledge on the cluster building blocks (hardware and software), peculiarities of cluster systems, their types, optimization and components. The difference between high performance, high availabillity, homogeneous and heterogeneous clusters is explained [2,15,16]. Finally, the recipes on cluster design and benchmarking are provided. At this point, students have enough knowledge to start a preparation of the report in the form of the proposal for specific cluster design [7]. This work requires from a student to "build" the best cluster indended for scientific tasks within an allocated budget and using current marketing information available on the Internet. The choice of the components and the interconnection type should be justified, and moreover, the students should think up of the real task, suitable for realization on the proposed cluster, and show its flow-diagram. Thus, as a results of this proposal preparation, the students obtain practical knowledge on cluster system planning and possible applications.

The course topic 7 (Table 1) introduces into the cluster "middleware", which is responsable for transparency, availability and scalability of the cluster system [17]. The connection between the middleware and the Single System Image (SSI) is described, and the SSI benefits, functions and realisation levels are given. As an example of SSI, the students are introduced into the openMosix technology [18]. The openMosix is a Linux kernel extension, which turns a network of ordinary computers into a clustering platform that is scalable and adaptive. The use of openMosix allows the cluster to adapt itself automatically to the workload, so that the resources allocation is continuously optimized. The cluster behaves much as does a Symmetric Multi-Processor (SMP), but this solution scales to well over a thousand nodes. It is important to note that there is no need to program applications specifically for openMosix, and nearly every Linux application automatically and transparently benefits from the distributed computing concept of openMosix. To demonstrate this facinating technology to students, two open-source LiveCD Linux distributions are used within the course [19,20].

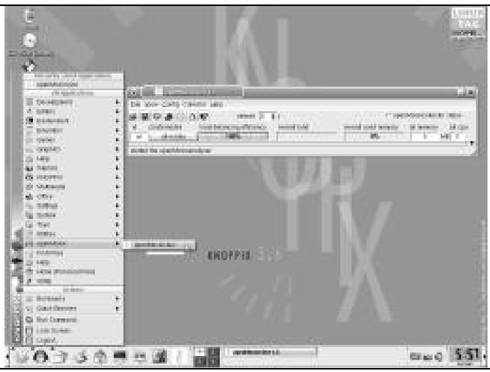

FIGURE 5. CLUSTERKNOPPIX LINUX LIVECD [19] WITH OPENMOSIX.

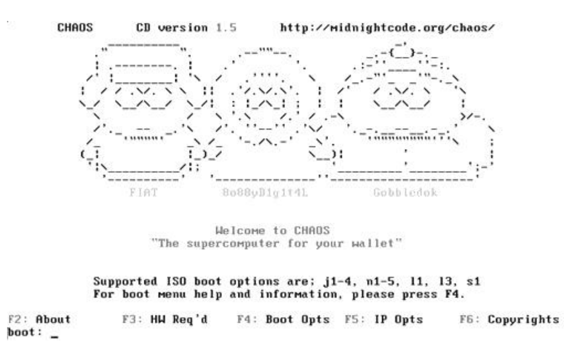

FIGURE 6. CHAOS Linux LiveCD [20] with openMosix.

The first one is the ClusterKnoppix Linux [19], which is based on Debian Linux with the openMosix kernel. The ClusterKnoppix distribution boots and runs entirely from CD and uses default KDE XWindows environment (Figure 5). Besides, this distribution includes hundreds of quality open source programs, which can be used for demonstration purposes. In particular, it has a graphic visualisation interface to openMosix, thus allowing to see and control cluster load in real time. Another Linux distribution, called CHAOS [20] (Figure 6), fits to a tiny CD disc (about 9 Mbyte) and boots any i586 class PC (that supports CD or PXE booting) into a working openMosix node, without disturbing the contents of any local hard disk. Once booted, CHAOS runs entirely from memory, that makes it an ideal system for class-room demonstrations. In fact, only about 10-15 minutes is required to convert all computer class into a cluster using CHAOS Linux distribution.

Since most clusters use UNIX-type operating system, the short introduction (topics 9 to 11) to Linux and basic knowledge on its installation, configuration and administration are provided within the course [21]. Students get information on Linux components, kernel, boot process, file system and processes (interactive, batch and demons).

The final parts of the course (topics 12 to 17) are related to the practical use of clusters. They introduce students into existing parallel programming technologies and main steps of parallel program design. Two main standards, the Message-Passing Interface (MPI) [3,8,22,23] and Parallel Virtual Machines (PVM) [4,8], are considered in details. Besides, the concept of the parametric execution system is discussed on the example of the „SPIDER" system [24], developed at the Institute of Solid State Physics of Latvian University by the author. The practical implementation of these technologies is illustrated on two examples: Monte Carlo modelling and graphical image rendering, which represent highly parallelizable problems.

3. CONCLUSIONS

A new introductory course on cluster computing has been developed, including supporting information, materials and technologies. It can be given using any available computer class-room facility and is based on open source type software, easily and freely obtainable from the Internet. The course has been introduced into the teaching process at Latvian University and successfully passed approbation process. Future development of the course can be performed in the direction of small dedicated cluster setup, that will provide indispensable support for student practical works.

ACKNOWLEDGEMENT

This work was supported in part by the Latvian Government Research Grants No. 01.0807 and 01.0821.

REFERENCES

[1] Kuzmin A. (2003) Cluster approach to high performance computing. *Computer Modelling and New Technologies* **7,** 7-15.

[2] Buyya R. (ed.) (1999) *High Performance Cluster Computing: Systems and Architectures.* Prentice Hall PTR.

[3] http://www-unix.mcs.anl.gov/mpi/

[4] http://www.epm.ornl.gov/pvm/pvm_home.html

[5] http://www.top500.org/

[6] http://www.cfi.lu.lv/lasc

[7] http://www.cfi.lu.lv/cluster

[8] Немнюгин С., Стасик О. (2002) *Параллельное программирование для многопроцессорных вычислительных систем*. БХВ, Санкт-Петербург.

[9] Шпаковский Г.И., Верхотуров А.Е., Серикова Н.В. (2004) *Руководство по работе на вычислительном кластере*. БГУ, Минск.

[10] Spector D. (2000) *Building Linux Clusters*. O'Reilly.

[11] Vrenios A. (2002) *Linux Cluster Architecture*. Sams.

[12] Sterling T. (2003) *Beowulf Cluster Computing with Windows (Scientific and Engineering Computation).* Wiley.

Proc. 3st Int. Conf. "Information Technologies and Management", April 14-15, 2005 (Information Systems Management Institute, Riga, Latvia, 2005) pp. - .

[13] Bookman C. (2002) *Linux Clustering: Building and Maintaining Linux Clusters*. Que.
[14] Kopper K. (2004) *The Linux Enterprise Cluster*. No Starch.
[15] Marcus E., Stern H. (2003) *Blueprints for High Availability*. 2nd edition. Wiley.
[16] Piedad F., Hawkins M. (2000) *High Availability: Design, Techniques and Processes*. Prentice Hall PTR.
[17] Buyya R. (ed.) (1999) *High Performance Cluster Computing: Programming and Applications*. Prentice Hall PTR.
[18] http://openmosix.sourceforge.net/
[19] http://bofh.be/clusterknoppix/
[20] http://itsecurity.mq.edu.au/chaos/
[21] Tansley D. (2000) *Linux and Unix Shell Programming*. Addison-Wesley.
[22] Pacheco P. (1996) *Parallel Programming With MPI*. Morgan Kaufmann Publishers.
[23] Шпаковский Г.И., Серикова Н.В. (2002) *Программирование для многопроцессорных систем в стандарте MPI*. БГУ, Минск.
[24] Kuzmin A. (2004) *Latvian SuperCluster - LASC: recent developments*. In Proc. 2nd Int. Conf. “Information technologies and management 2004“, Information Systems Management Institute, Riga, Latvia.