AC 2011-1301: IMPLEMENTING AN AFFORDABLE HIGH PERFORMANCE COMPUTING PLATFORM FOR TEACHING-ORIENTED COMPUTER SCIENCE CURRICULUM

Jeongkyu Lee, University of Bridgeport

Jeongkyu Lee received a B.S. from Sungkyunkwan University in Mathematic Education and an M.S. from Sogang University in Computer Science, both of Seoul, Korea in 1996 and 2001, respectively. Before he pursued his doctorate, he worked as a database administrator for seven years with companies including IBM. In fall 2002, he entered the Doctoral program in Computer Science and Engineering at the University of Texas at Arlington. After he received Ph.D. degree in summer 2006, he joined the Department of Computer Sciences and Engineering at University of Bridgeport, CT as an assistant professor. His primary research interest is in the multimedia database management. His work also includes techniques for multimedia data mining, video processing, multimedia ontology, and medical imaging. He is a member of IEEE, ACM, and ASEE.

Omar Abuzaghleh, University of Bridgeport

Mr. Abuzaghleh is the Assistant Lab Manager in the School of Engineering at the University of Bridgeport. He has research interests in the areas of distributed database systems, data mining, and cloud computing. Mr. Abuzaghleh received the B.Sc. degree in computer science and applications from Hashemite University in 2004, the MS Degree in computer science from university of Bridgeport in 2007. Mr. Abuzaghleh is currently pursuing his Ph.D. degree in computer science and engineering at university of Bridgeport.

Implementing an Affordable High Performance Computing Platform for Teaching-oriented Computer Science Curriculum

Abstract:

With the advances in computing power, High Performance Computing (HPC) platform has impact on not only scientific research in advanced organizations but also computer science curriculum in the educational community. For example, multi-core programming and distributed and parallel system are highly desired courses in the computer science classroom. However, the high cost of HPC equipments and maintenance makes it difficult to be employed in conventional computer science curriculum. Specifically, teaching-oriented institutions cannot afford the HPC due to the high cost, lack of experience and less research infrastructure.

The main objective of this paper is to present an affordable and easy-to-use high performance cluster system that can be used in teaching-oriented computer science curriculum. In order to address this, we design and implement an affordable high performance cluster system based on PlayStation 3®. PS3 is a well-known game console manufactured by Sony. Since each PS3 console has a IBM Cell BE processor that consists of 8 Synergistic Processing Elements (SPEs) and 1 Power Processing Element (PPE), it can be used as a processing node with multiple-core processor in the cluster system.

As a result, the PS3 Cluster system provides enough computing power as a HPC for computer science courses, while the total cost is less than 10% of existing cluster system in market that has the similar performance. In addition, the implemented cluster system has been used for new and existing computer science courses, such as CPSC 592: Parallel and Distributed Database, CPSC 590: Parallel and Distributed Processing, and CPSC 591: Parallel Programming.

1. INTRODUCTION

In the era of internet with ever-growing information and data, it is highly desired for the industry and academy to utilize high computational power of High Performance Computing (HPC) in their business and educational activities. Specifically, cluster-based systems dominate the market of HPC, since they are more flexible and scalable comparing to the computing power of single-node super computer. According to the list of top 500 super computer sites[1], the cluster-based super computer has increased more than 80% since 2000. Therefore, there exists a great need of skilled and educated students and engineers who are able to design, implement and operate such HPC environments. For example, the students in computer science and engineering aspire to have multi-core programming and distributed and parallel system. However, the high cost of HPC equipments and maintenance makes it difficult to employ them into the conventional computer science curriculum. Specifically, teaching-oriented institutions cannot afford HPC platforms due to the high cost, lack of experience and less research infrastructure of it.

In order to offer HPC related courses such as Distributed and Parallel system, Parallel Programming or Distributed and Parallel Database, in teaching-oriented institutions, we need affordable infrastructure where student and instructor can design, implement and test the HPC

system. An alternative approach to the HPC platform is the cluster-based computing system for its flexibility and feasibility of high computing power with relatively low cost. In addition, open source software including operating systems (i.e., Linux) and programming SDK help to reduce the cost, even make us easy to implement and use them.

Therefore, we design and implement an affordable high performance cluster system on the PlayStation 3®. PS3 is a well-known game console. Since each PS3 device has a IBM Cell BE processor that consists of 8 Synergistic Processing Elements (SPEs) and 1 Power Processing Element (PPE), it can be used as a processing node with multiple-core processor in the cluster system. 6 PS3 are connected with each other through gigabit ethernet to form a cluster system. Each node is running on Fedora Core 12 Linux and the parallel programming software are IBM Cell SDK and Open-MPI. All of them are open source software, hence they can be used easily in the classroom. This lineup is not only for the educational purpose, but can be used as military and industrial solutions.

In this paper, we present an affordable and easy-to-use high performance cluster system that can be used in the classroom of a teaching-oriented computer science curriculum. This paper is organized in the following way. In Section 2, related works in academic HPC platforms are discussed. Section 3, we describe our design and implementation of the PS3 based cluster system. In addition, newly developed courses are presented in Section 4. Finally, concluding remarks and future works are discussed.

2. RELATED WORKS

There exist wide variety of courses offered at universities in the US regarding to high performance computing and they can be classified into the following two groups[8,9]:

- *System Architecture*, where the aim is to teach students how to design, to implement and to deploy a HPC platform. Incomplete list of related course include: Theory of Parallel System, Introduction to Distributed System, High Performance Parallel Computing, Parallel and Distributed Computing.
- *Application*, where the aim is to implement high performance distributed and parallel program on the HPC platform. These courses could be Parallel and Distributed Programming, Distributed and Parallel Database, Multi-core Programming, and Performance Evaluation.

Table 1 is a survey of high performance computing courses offered in the US only. As shown in the table, most HPC courses utilize high performance computing infrastructure such as supercomputer or cluster systems. For example, the Berkeley Network of Workstations (NOW) project builds a distributed supercomputer using a network of workstations, which is mainly sponsored by NSF[10]. Currently, NOW includes 100 SUN UltraSPARCs and 40 SUN SPARCs running on Solaris, 35 Intel PC's running on Windows NT or on PC UNIX variant. All are connected by a Myrinet switched network. The NOW project stimulated many HPC related courses, such as CS 258 Parallel Processors. However, since the universities listed in Table 1 are research intensive schools, the supercomputers or cluster systems are mainly sponsored by the government agency, such as NSF or NIH, or the industry.

Although there exists a great needs of HPC courses in teaching oriented schools, it is very hard to build such HPC system for the high cost of super stations and the lack of experience in HPC.

Table 1. Survey of HPC courses in US

University	Course Name	Facility	# of Cores (or # of Node)	Sponsor	Link
University of California, Berkeley	CS 258 Parallel Processors	NOW	175 nodes	NSF	http://now.cs.berkeley.edu/
University of California, Berkeley	CS267 Applications of Parallel Computers	NERCE	153408 nodes	DOE	http://www.nersc.gov/
MIT	Applied parallel Computing	BEOWULF	16 nodes (Intel Xeon)		http://beowulf.lcs.mit.edu/
University of Utah	CS6963: Parallel Programming for GPUs	GPU	3 units		http://www.cs.utah.edu/~mhall/
University of California, San Diego	CSE225 Grids and High Performance Distributed Computing	FWGrid	100 nodes	NSF	http://fwgrid.ucsd.edu/
Penn State University	CS437 Introduction to Parallel Computing	CyberStar	2048 cores		http://rcc.its.psu.edu
University of North Carolina	COMP 633 Parallel Computing	BASS	452 CPU	NIH	http://wwwx.cs.unc.edu/Research/bass
John Hopkins University	CS420: Parallel Programming	HHPC	1200 Intel cores		http://idies.jhu.edu/hpc.aspx
University of Maryland	CMSC 838T: Advanced Topics in Programming Languages	Red	16 dual linux		http://www.cs.umd.edu/class/spring2003/cmsc838t/

3. DESIGN AND IMPLEMENTATOIN OF PS3 CLUSTER

An affordable HPC platform needs to satisfy the following requirements: (1) *price*: the system should be afforded to a teaching intensive school in the price range between $1,500 to $2,000 US dollars, which is the basic amount of money supported for newly offered course at the University of Bridgeport, (2) *processor power*: the processor power is one of important elements to determine if the system is HPC compatible. Therefore, the processor of the HPC node needs to be multi-cores, (3) *scalability*: the system should be flexible to scale up in order to support parallel processing, and (4) *easy to use*: due to the lack of experience of high performance computing in the department, the system should be easy to use and maintain.

In order to satisfy the aforementioned requirements, we adapt PlayStation 3® (PS3) based cluster system for the HPC environment. PlayStation, PS in abbreviated, is a video game console manufactured by Sony. Since the first PS was released in 1994, there have been three main generations, i.e., PlayStation, PlayStation 2 in 2000, and PlayStation 3 in 2006[2]. After PS3 was

announced, the console has attracted considerable attention as a computing node because of the high performance of CELL processor in the PS3 and its significantly low price while comparing to other processors. CELL processor is jointly developed by the alliance of Sony, Toshiba, and IBM, called STI in 2005. On 2007, the first PlayStation 3 based academic computing cluster was created at North Carolina State University. Prof. Frank Mueller used 8 PlayStation 3 to create the computing cluster[3].

In Spring 2010, Computer Science and Engineering of University of Bridgeport offered a new course, CPSC592 Distributed and Parallel Database, where a PS3 based cluster system, called UB PS3Cluster, was designed and implemented in that semester. Six PS3 are used to form a cluster via Gigabit Ethernet. Figure 1 shows the overview of UB PS3Cluster implemented. Basically, the cluster consists of six PS3 storage as computing nodes, a gigabit ethernet for communication, and a front node PC.

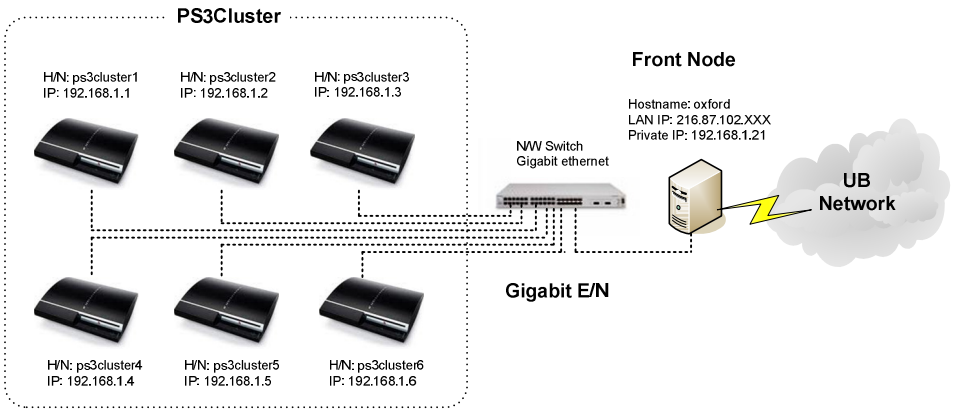

Figure 1 Overview of PS3 Cluster with 6 PlayStation 3

Hardware

A PS3 has many interesting features, for example, it supports gigabit ethernet, a blue-ray, a graphic card, and other OS. However, the most valuable one is the IBM Cell BE processor. A Cell processor consists of a dual-threaded, 64 bit power architecture compliant Power Processor Element (PPE), and eight Synergistic Processor Elements (SPEs). These units are interconnected by an Element Interconnect Bus (EIB) that supports a bandwidth of 25.6 gigabits per second. In other words, each PS3 console can support eight parallel processing using 8 SPEs[4]. A Cell processor used in PS3 utilizes 6 SPEs for working elements, since two of them are reserved for the game OS. In addition, PS3 console has 256 MB of XDR DRAM, gigabit ethernet network card, and NVIDIA graphic card with 256 MB video RAM. Therefore, PS3 can be treated as a computing node in a cluster computer of a HPC platform. The hardware specification used in this project is defined as follows:

(1) PlayStation 3 (6 nodes)

- CPU: IBM Cell BE processor
    - 1 PPE (Power Processing Element)
        - 64-bit PowerPC

        - 32 KB Instr/Data L1 Cache, 512 KB L2 Cache
        - SMP (2 threads)
        - 3.2 GHz (SP: 25.6 GFlops DP: 6.4 GFlops)
    - 8 SPEs (Synergistic Processing Element)
        - 256 KB Local Store
        - 3.2 GHz (SP: 25.6 GFlops DP: 1.83 GFlops)
        - 128-bit Vector Registers
    - EIB (Element Interconnect Bus)
        - Interconnects PPE, SPEs, Memory, I/O
        - Simultaneous Read/Write
    - MIC (Memory Interface Controller)
        - Interfaces to XDR Memory
        - Theoretical B/W of 25.6 GB/s
    - 218 GFLOPS of total floating point performance
    - 25.6GB/sec memory bandwidth
- Memory: 256 MB of XDR DRAM
- Network Card: GigaBit Ethernet network card
- Graphic Card: NVIDIA and 256 MB of video RAM
- Hard Disk: User-upgradeable 2.5" SATA hard drives, 80 GB
- Blue-ray Disc drive, 2 X USB 2.0, HDMI 1.3a, and Wi-Fi

(2) Gigabit Ethernet Switch

- CISCO 8PT 10/100/1000 GBIT SWCH

(3) Front Node

- HP Pavilion PC with 2 NIC (one for UB network and another for private network of cluster)
- DHCP server for private network
- Cluster can be accessed via Front node

Software

PS3 console allows the installation of other operating system in addition to the PS3 Game OS. Therefore, Fedora Core 12 Linux is installed on each node of the UB PS3cluster. Fedora Core 12 is running on top of a virtualization layer of the Game OS; therefore, only limited operating systems can run on PS3. The available OS for PS3 include Fedora Core, Yellow Dog, Gentoo Power PC 64 edition and Debian. In order to support clustering programming, Message Passing Interface (MPI) standard is employed and Open MPI (version 1.4.2) is installed in the system because of its high performance message passing library. In addition, IBM Cell SDK 3.0 provides the development environment for programming the CELL processor with GNU gcc complier.

(1) Operating System: Fedora Core 12 with 2.6.32.11-99.fc12.ppc64 Linux kernel

- Fedora Core 12 with 2.6.32.11-99.fc12.ppc64 Linux kernel

(2) Open MPI

- High Performance Message Passing Library
- Version: 1.4.2
- Website: http://www.open-mpi.org/

- Compiler: mpicc
- Execute: mpirun

(3) Cell SDK 3.0

- CELL SDK provides the development environment for programming the CELL processor
- Version: 3.0
- Website: http://www.ibm.com/developerworks/power/cell/
- Compiling
    - Compile SPE using spu-gcc
    - Embed it to a specific variable using ppu32-embedspu
    - Compile PPE and link to SPE using gcc

(4) Eclipse

- IDE with tooling for various languages:
- Version: 3.5.1
- Website: http://fedoraproject.org/wiki/Eclipse

Implementation

(1) Building PS3 Cell Cluster

Six PS3 (80 GB, original model, released in 2008) are used for the cluster because the newer version of PS3 (120 or higher GB, slim model) removes the Other OS feature from firmware version 3.21. To build a cluster system, they are connected to a gigabit ethernet switch. In addition, a front node PC is attached to the system via the switch for security purpose. The front node is used for login, security, monitoring and management purpose (including console interface) of the UB PS3cluster. It has two network interface cards (NICs); one connected to the public network while the other is connected to the private network of the clusters. Figure 2 shows some pictures of implemented PS3 cluster.

Figure 2 Pictures for implemented PS3 Cluster

(2) Installing Fedora Core OS

As mentioned above, the PS3 Game OS, up to firmware version 3.21, allows the installation of other operating systems on top of its virtualization layer, hypervisor. Fedora Core 12 Linux is installed as following steps:

- Install Fedora Core 12 Bootloader, i.e., petiboot loaer
- Install Fedora Core 12 Linux using DVD
- Configure network interface
- Install base packages

- Install packages and software required in cluster, i.e., NFS, SSH, DHCP server and client, and X server
- Configure users and groups

(3) Installing Cell SDK and Open MPI
After installing operating system, Cell SDK 3.0 and Open MPI are installed for software development and cluster programming, respectively.

4. COURSE DEVELOPEMENT

UB PS3cluster is designed and implemented in CPSC 592: Parallel and Distributed Database offered by Computer Science and Engineering department in Spring 2010. A course syllabus consists of two parts: Distributed Database Systems and Parallel Database systems. One objective of this course is to provide the fundamental and advanced concepts and techniques of parallel and distributed database to fulfill the needs in both academia and the industry. The other objective is to prepare the student for furture research in the area of very large database systems. During the semester, we implement Distributed and Parallel DBMS using UB PS3cluster.

A list of DBMSs that support parallel and/or distributed database management system as outlined in Table 2.

Table 2. Summary of Parallel and Distributed DBMS

Parallel and Distributed DBMS	Architecture	Query	Partition	Fedora Linux
DB2 Parallel Edition[5]	Shared Nothing	Cost-based parallel optimizer	Table Partitioning	No
Informix Extended Parallel Server[6]	Shared Nothing	Parallel query	Data/Index Partitioning	No
Oracle Real Application Cluster	Shared Everything	Parallel query management	Table/Index Partitioning	No
MySQL Cluster[7]	Shared Nothing	Parallel plan	User defined partitioning	Yes

As shown in Table 2, MySQL cluster is the only Parallel and Distributed DBMS that support Fedora Core Operating System. Therefore, we installed MySQL cluster 7.0.13 on Fedora Core 12 Linux. In addition, since MySQL Cluster employs s*hared nothing* architecture, each node owns its memory and storage area for the instance. MySQL cluster has three different nodes: (1) *management nodes*: manage the entire distributed database in the cluster, (2) *SQL nodes*: coordinate SQL query requests, and (3) *data nodes*: have partitioned tables and indexes.

Figure 3 is an overview of installed MySQL Cluster on UB PS3cluster and its configuration. It consists of 1 management node, 2 SQL nodes and 3 Data nodes. However, they can be easily scaled up by adding PS3 node to the cluster.

During the semester, the student of CPSC592 conducted the following class activities using PS3cluster:

- Lectures (8 weeks)
- Open seminars (2 weeks)
- Design and implementation of PS3cluster (3 weeks)
- Installation and configuration of MySQL cluster DBMS (3 weeks)
- In addition, two projects are conducted: (1) modeling partitioned database, and (2) implementing Distributed database using UB PS3cluster.

In addition to the Distributed and Parallel Database course, the following courses are possible to be offered based on UB PS3cluster:

- CPSC 590: Parallel and Distributed Processing, and
- CPSC 592: Parallel Programming.

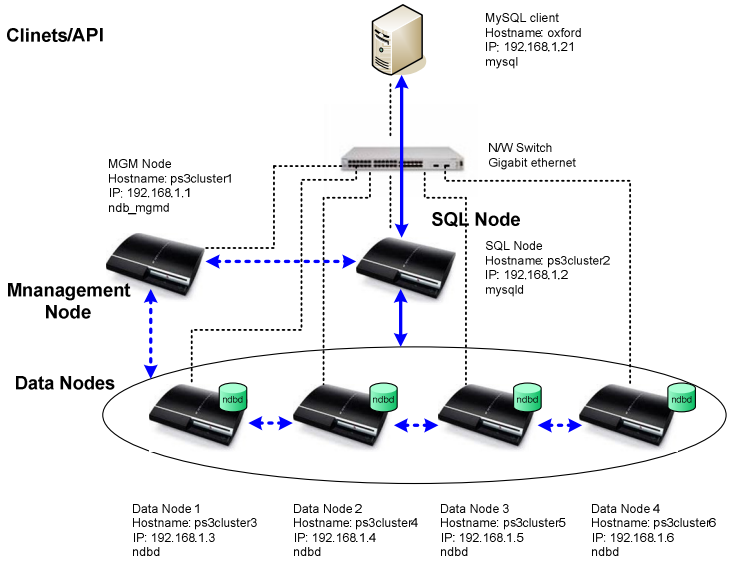

Figure 3 MySQL Cluster (DDBMS) installed on UB PS3 Cluster

5. CONCLUSION

In this paper, we presented our implementation of an affordable high performance computing (HPC) platform for teaching-oriented computer science program at University of Bridgeport. Although HPC is highly desirable for undergraduate and graduate students, the high cost of HPC equipments and maintenance makes it difficult to adapt them into the conventional computer science curriculum. Specifically, teaching-oriented institutions cannot afford the HPC platform due to the high cost, lack of experience and less research infrastructure of it. In order to address this, we have designed and implemented an affordable high performance cluster system based on PlayStation 3®(PS3). PS3 is a well-known for game console. Since each PS3 has a IBM Cell BE processor that consists of 8 SPEs (Synergistic Processing Elements) and 1 PPE (Power Processing Element), it can be used as a multiple-core processing node of the cluster system. As a result, the PS3 cluster provides enough computing power as a HPC platform, costing less than 10% of the set up cost of existing cluster system in market with the similar performance. The implemented clustering system was used in both new and existing computer science courses. In the future, more

HPC related courses will be offered, e.g., CPSC 590: Parallel and Distributed Processing and CPSC 591: Parallel Programming.

REFERENCES

1. Top 500 Supercomputer Sites, http://www.top500.org/
2. Sony Business Development, http://www.scei.co.jp/corporate/data/bizdatajpn_e.html
3. Sony PS3 Cluster (IBM Cell BE), http://moss.csc.ncsu.edu/~mueller/cluster/ps3/
4. J. A. Kahle, M. N. Day, H. P. Hofstee, C. R. Johns, T. R. Maeurer, and D. Shippy. "Introduction to the cell multiprocessor", IBM J. Res. Dev. 49, 4/5 (July 2005), 589-604.
5. Chaitanya Baru and Gilles Fecteau. "An overview of DB2 parallel edition", In Proceedings of the 1995 ACM SIGMOD international conference on Management of data (SIGMOD '95), New York, NY, 1995
6. Andreas Weininger. 2000. "Handling very large databases with informix extended parallel server", SIGMOD Rec. 29, 2 (May 2000), 548-549.
7. Sudsanguan Ngamsuriyaroj and Rangsan Pornpattana. "Performance Evaluation of TPC-H Queries on MySQL Cluster". In Proceedings of the 2010 IEEE 24th International Conference on Advanced Information Networking and Applications Workshops (WAINA '10). IEEE Computer Society, Washington, DC, USA.
8. Thomas Hacker, "Hands-on High Performance Computing: Developing a Cluster Computing Course for Real World Supercomputing". Proc. of the 2010 National ASEE conference, Vancouver, BC, 2010.
9. Afsaneh Minaie, Reza Sanati-Mehrizy, "Incorporating Parallel Computing in the Undergraduate Computer Science Curriculum", Proc. of the 2009 National ASEE conference, Austin, TX, 2009.
10. David E. Culler, "Parallel computing on the Berkeley NOW", In Proceedings of the 9th Joint Symposium on Parallel Processing (JSPP 97), 1997.