Cluster Computing in a Computer Major in a College of Criminal Justice*

Douglas E. Salane
dsalane@jjay.cuny.edu

and

Boris Bondarenko
websitedeveloper@aol.com

John Jay College of Criminal Justice
Mathematics Department
December 23, 2003

Extended Abstract
Paper to be presented at the

2004 USENIX Annual Technical Conference
June 27–July 2, 2004, Boston, Massachusetts

---

* This work supported in part by a NASA Curriculum Improvement Partnership Award

Introduction The computer major at John Jay College is a systems oriented curriculum. The program focuses on computing and mathematical methods that improve the practice of criminal justice and public agency functions. Computer clusters are of interest to these agencies because they offer high throughput, highly available platforms for mission critical applications. In addition, they offer a low cost platform for the high performance computing required in simulation and security research.

The computer program at John Jay College uses cluster computing in two ways: 1) to enhance the learning experience in computer courses, and 2) to support research in areas of interest to law enforcement and public agencies. Since cluster computing brings together the latest developments in operating systems, file systems, high performance networking, algorithms, computational mathematics and computer architecture, a computer cluster enables students to gain hands-on experience exploring the latest developments in many different areas of computing.

This talk will review several on going projects that use cluster computing to support educational and research activities. The talk will discuss the type of cluster we have built and why it is appropriate for an institution that has limited support personnel for cluster management. Finally, the talk will outline several cluster developments that we think will help promote the use of computer clusters in smaller institutions with limited resources.

Using a cluster effectively when there is limited budget for hardware and personnel presents a number of challenges. Issues such as knowledgeable faculty and trained support staff, building and maintenance of cluster computing facilities, adequate preparation of students in lower level courses, and selection of appropriate course material is an active area of discussion in the CS educational literature [1]. Fortunately, the Mathematics Department has a NASA Curriculum Partnership Improvement Award that has helped us acquire the necessary staff expertise, develop curricula, retrain faculty, and address many of the challenges. Details of the project can be found at web.math.jjay.cuny.edu.

The Department has constructed three clusters to support course work and research. These include a traditional Beowulf cluster for high performance computation, instruction in parallel computation and distributed algorithms; a database cluster to support work in large-scale databases and highly available, high throughput computing; and a small cluster to support student interactive Secure Shell (SSH) sessions from on campus and home. With these facilities, we are able to introduce future IT professionals to the benefits of cluster computing and support research in large-scale database systems and high performance computing. The following are some specific on going educational and research activities that we will discuss.

Parallel programming environments Parallel computing in a message-passing environment is covered in our advanced computer networking and operating systems courses. Using the computational cluster and a performance analysis program called Netpipe [6], students have been able to perform a detailed analysis of the MPICH send/receive semantics [4]. They have duplicated the results in [6] that show a drop off in performance as message sizes grow beyond the system allocated buffer sizes.

High performance matrix computations Large-scale matrix computations are required to solve the underlying mathematical problems in modeling and simulation. In addition, singular value decompositions are used in the principal component analysis of large data sets. The computational cluster provides the ScaLAPACK Library (www.netlib.org/scalapack/scalapack-home.html) for solving systems of equations, calculating eigenvalues and singular value decompositions, and solving least square problems.

Critical to the performance of a ScaLAPACK is an effective implementation of the Basic Linear Algebra Subroutines (BLAS) [3]. Performance results for a key routine that performs a block matrix multiply will be examined for the three of the most widely used BLAS versions: the Intel Math Kernel Library version (www.intel.com), ATLAS generated BLAS (http://math-atlas.sourceforge.net), and the KGoto version (www.utexas.edu/users/kgoto). The effect of cache and TLB misses on the performance of each implementation will be discussed briefly. The performance of each in the High Performance Linpack benchmark (www.netlib.org) will be presented.

Distributed Algorithms A cluster is an excellent environment for the research and course work in distributed algorithms. Since cluster nodes trust one another, a distributed implementation can spawn a task on each node. Unlike a simulation environment, students and researchers can study synchronization issues that arise in a live implementation of a distributed algorithm. Using UDP, students in our networking courses develop a Perl implementation of the Distributed Bellman Ford algorithm that runs in our cluster. An x-windows interface we constructed allows students to view the routing tables as they are built on each node.

Password cracking Clusters offer the disk storage and computational power required for computing and storing encrypted passwords and then using them in a password cracking attacks. Such computational resources were heretofore limited to groups with extensive computational resources. Increasingly, law enforcement agencies are concerned about the types of attacks possible using "modest" high performance computing resources. Password cracking results on our cluster using software developed at San Diego Super Computer Center [5] will be presented.

Highly available databases and on-line services We have two on-going projects to investigate the use of clusters for high throughput and high availability. The first is a database project. The FBI's National Incident Based Reporting System (NIBRS) provides an annual list of crimes reported each year (http://www.fbi.gov/ucr/ucr.htm). The department has built a relational database that houses over 10 years of NIBRS data. We are building a portal to make this data available to the criminal justice community. We are using a Linux Virtual Server to load balance between two computers that house the NIBRS database. In a second project, we are using the Linux IPTables router and a simple routing technique to load balance two computers that handle interactive SSH sessions.

Management Considerations Following the original Beowulf developers, we have found it is essential to stick to well-supported, commodity hardware and open source or freely available software. Besides pushing us over budget, we have found that most proprietary solutions, which often offer the highest levels of performance, are difficult for us to deploy and manage. Our 12 node, 24 CPU Linux based computational cluster includes a gigabit Ethernet interconnect, an NSF file system, and a separate world node that provides access but is not used for computation. Since cluster nodes become obsolete in about a year and a half, our cluster is a pile of PCs rather than a blade-based cluster. Unlike blades, old cluster PCs find other uses after life in a cluster. In order to effectively use and manage a cluster, we have found that it is essential that in-house staff develop expertise in all required cluster technologies, including extensive Linux expertise, expertise in middleware such as MPICH, and in applications such as ScaLAPACK.

Concluding Remarks A cluster is a complex computer system in which many layers of software and hardware must function correctly if the system is to produce valid results. The cluster community must continue to develop the software and hardware technologies that will allow commodity clusters to offer a single system image [2] and high performance. Clusters must be easy to use and manage. Moreover, it must be easy to detect and correct hardware problems. Too often we have relied on applications level software, e.g., the BLAS, to detect subtle errors in the cluster communication hardware and software. Cluster installation software such as ROCKS (rocks.npaci.edu) and OSCAR (sourceforge.net/projects/oscar) facilitate cluster implementation and address the OS version skew problem. However, problems with integration of applications, middle-ware, operating systems and C libraries still force end cluster users like us to be intimately familiar with what goes on under the hood. Finally, publicly available, inexpensive, well-supported, commodity, high performance interconnects and cluster file systems must be developed to allow smaller institutions to reap the full benefits of cluster computing.

References

1. Apon, A., Rajkumar, B., Jin, H., Mache, J., Cluster Computing in the Classroom: Topics, Guidelines, and Experiences, Proceedings of the First International Symposium on Cluster Computing and the Grid (CCGRID '01), IEEE Computer Society, August 2001.

2. Buyya, R., High Performance Computer Clusters: Archtecture and Systems, Vol I, Prentice Hall, Englwood Cliffs, NJ, 1999.

3. Dongarra, J.,, Du Croz, J.,, Hammarling, S.,, Duff, I.,, A Set of Level 3 Basic Linear Algebra Subprograms, ACM Transactions on Mathematics Software, Vol. 16, No. 1, pp. 1-17, March 1990.

4. Gropp, W., Lusk, E., Doss, N. and Skjellum, A., A high-performance, portable implementation of the MPI message passing interface standard, Parallel Computing, Vol. 22, No. 6, pp. 789-828, 1996.

5. Perrine, T. and Devin, K., Teracrack: Password Cracking Using Teraflop and Petabyte Resources, SD2003,available at http://security.sdsc.edu/publications/terracrack.pdf

6. Turner, D., Chen, X., Protocol-Dependent Message Passing Performance on Linux Cluster, 2002 IEEE International Conference on Cluster Computing, IEEE Computer Society,Chicago, IL, September 23, 2002, pp. 187-194