Prototyping an Object Warehouse

T. Sasaki[1], *H. Okazawa*[2], *A. Shibata*[1], *N. Takashimizu*[3] *and Y. Watase*[1]

[1] Computing Research Center, KEK
[2] Shizuoka Seika College
[3] Department of Mathematics and Computer Science, Shimane University

Abstract

We present the design and status of a prototype object warehouse, an OO version of a Data Warehouse. Data analysis in HEP is the process of discovering meaningful new correlations, patterns and trends by scanning through large amounts of stored data, using statistical and mathematical techniques as well as pattern recognition technologies. This method is called as ”data mining” in the information science. A large information center specialized for fast data mining, a Data Warehouse, has been implemented in many organizations. We have fully taken knowledge in these areas into account designing the system.

We are also developing software tools for an Object Warehouse based on already existing data mining techniques for fast and efficient data analysis in HEP.

Keywords: ODBMS, SAN, data mining, data warehouse

1 Introduction

> *“Data mining is the process of discovering meaningful new correlations, patterns and trends by sifting through large amounts of data stored in repositories, using pattern recognition technologies as well as statistical and mathematical techniques.” Gartner Group*

What people doing in the HEP data analysis is nothing special, but similar things have done in many other fields sometimes with much sophisticated and more systematic way. The process of discovering meaningful new correlations, patterns and trends by scanning through large amounts of stored data, which exactly we do in our analysis, using statistical and mathematical techniques as well as pattern recognition technologies, is called as ”data mining”. In conjunction with knowledge from artificial intelligence and machine learning research field, similar technologies are often called as KDD (Knowledge Discovery in Database). In this paper, we use the term, “data mining”, to focus on data analysis, but this does not mean that we hesitate to use KDD methods such as Bayesian networks which will be discussed in the section 2.

Data mining is already very popular in many fields. The main reason why these techniques are never used in HEP is simply that we stored data sequentially on tapes. To handle multi-dimensional relationship, database is essential for data mining. Advances of storage and microprocessor technologies, with new algorithms, give us the opportunity of introducing this technology in HEP. Now, we are not necessary to hesitate to use these existing proven technologies, because many of experiments have decided to use ODBMS for their analysis.

To do datamining efficiently, it is necessary to have a big information center, data warehouse, where one can handle huge amount of multidimensional data quickly. We built a prototype for OO version of the Data Warehouse, the Object Warehouse, based SAN (Storage Area Network) technology. We present the current status in the section 4.

2 Data Analysis Methods

In conventional HEP analysis, much iteration is necessary to decide event selection criteria, which is sometimes biased from theoretical predictions. This is happen because we treated high dimensional relationship with complicated correlation, without having proper statistical and mathematical methods. Neural networks are only method used widely in HEP, *e.g.* in B tagging. However, neural networks have problems such like, over learning, difficulty of understanding the relation between final result and input variables because of existence of intermediate layers, difficulty of estimation of systematic errors and so on.

We will be able to analyze data much quicker with smaller computing resource if we can decrease the number of iteration. To answer this demand, we are developing software tools implementing new algorithms such as Bayesian networks, Markov Chain Monte Carlo and so on, to improve efficiency of analysis and also Monte Carlo simulation. A Bayesian network is a graphical model for probabilistic relationship among a set of variables [3]. Bayesian networks are already widely used for clustering statistical events in other fields. Bayesian statistical methods in conjunction with Bayesian networks will provide more efficient approach for avoiding overfitting of data in HEP also. They will be useful to select a particular type of events from data. Bayesian networks can be applied to the problem with and without supervised learning.

The AutoClass[3][5] Project at NASA developed the algorithms and software for unsupervised Bayesian networks and has great success in the analysis of IRAS(Infrared Astronomical Satellite) data[6]. This software is also applicable for image processing on Cherenkov counters and so on, today.

3 Object Warehouse

The size of database has already reached to tera bytes in also commercial systems, not only in scientific systems. Mostly RDBMS was used at such Date Warehouses to store data and handle high dimensional relationship. However, RDMBS has a technical limit on sizes and complexity of relationship. To solve the situation, ODBMS was thought to be useful[2]. However, just because the companies who make ODBMS are small, many of users employed ORDBMS instead. Due to rapid grow of data size, and higher complexity of relationship among items are very high, still there might be a chance to get ODBMS popularity in the future. From the same reasons, HEP experiments choose ODBMS, particularly Objectivity/DB.

We will attempt to store all of data in the format of *"persistent objects"* and the relationship among them on ODBMS, and apply our new software tools on the database. In this sense, our system is not a simple Data Warehouse, but an Object Warehouse. ODBMS provides good solution to handle high dimensional complex relationship among many objects.

To do data analysis quickly, it's necessary to process events in parallel. Recent advance in microprocessor technology makes computers very cheap and we have opportunity to build large-scale computer farm based on PC's easily. However, the processing speed concerning ODBMS might be limited by I/O or communication performance. In designing the system, we need to carefully think about these points to obtain enough performance on random access from many hosts.

4 Prototyping

As we described in the previous section, the key issue in data mining is the performance of databases accessed from many clients. Objectivity/DB, which is widely used in HEP, has a bottleneck at the I/O server, AMS. We improved the situation using a SAN (Storage Area Network) technology to share a disk among clients over the fiber channel.

We introduced the SANergy product made by Mercury [1] for file sharing among Fiber Channel attached computers. This product utilizes NFS mechanism for access locking. The MDC (Meta Data Controller) host looks like a NFS server from clients and files are accessible as the files are under a NFS mounted directory. However, real I/O is done through Fiber Channel. All of the commands are not necessary to be modified to access files under SANergy volumes transparently. Because Objectivity/DB uses usual files to store database, AMS is not necessary running to the disk-attached host and one can just use NFS also. Because of this feature, we could build our prototype system based on SAN very easily.

For the prototype system, we introduced one Windows NT machine as the MDC controller and 4 Solaris machines for clients as shown in the figure 1. SANergy is currently available on NT, Solaris, IRIX, AIX and MacOS.

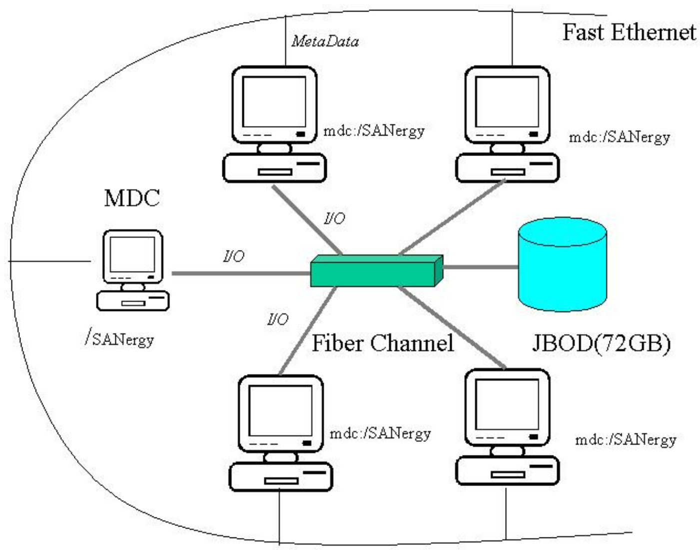

Figure 1: Prototype system

At first, we tried sequential I/O benchmark test on the system. It was shown that this system has good performance even with very short record size. For example, at 512 bytes record size, I/O rate is more than 30MBytes/sec on a client. Also for longer record length, the measured performance is almost flat and shows good behavior.

The OO7[7][8] benchmark program are widely used to measure ODBMS performance and measured results for many ODBMS systems are available. The original version of OO7 for Objectivity/DB was made for Objectivity version 2.X. We modified the code to run on Objectivity version 5.2. Using this benchmark program, we have measured the total performance of transactions versus number of clients simultaneously accessing. The detail of results will be presented somewhere else in the future. We present here the performance in single client access case only.

Comparing with the “gendb” result on the local case, the SANergy case was approximately 60 percent slow despite much higher performance in sequential I/O. Also the SANergy bench was slower even than NFS. The reason is thought to be that SANergy uses NFS server on a Windows NT machine for meta data management and access locking. Because the default page size of Objectivity/DB is 8KB and there are very frequent disk access, overhead of meta data management and locking might be accumulated in rather large numbers. Also NFS server performance on Windows

[1] Recently, IBM Tivoli Systems Inc acquired shared storage division of Mercury.

Table I: OO7 benchmark results

configuration	gendb med9	bench med9
SCSI III local disk	125	88
NFS Ver 3 (100BASE-T) mounted from another Solaris machine	422	131
SANergy volume	200	175

NT is thought to be slower than one on Solaris. After we will evaluate Solaris version of MDC, which will be released soon, we will investigate the reason and try performance tests again. However, the best solution seems to be that Objectivity provides the SAN version of AMS cooperated with Mercury if necessary. This does not seem difficult.

We are also planing to evaluate another product for file sharing on SAN from Fujitsu. We will compare performance of SANergy and the other product to do further research on the storage system for ODBMS.

5 Summary

Modern data mining technology was found to be also useful in HEP. We are developing the software tools using new algorithms for event analysis and also Monte Carlo simulation to decrease the number of unnecessary iteration. To use these tools, an Object Warehouse is required. We have build the prototype and tested the performance.

Toward the ATLAS Japan regional center, we will continue to do further research on the software tools, new algorithms, database technologies and also storage technologies to provide a competitive analysis environment.

References

1 Usama M. Fayyad *et. al.*, “Advance in knowledge discovery and data mining”, The MIT Press, 1996, ISBN:0-262-56097-6.

2 Neilson T. Debevioise, “The Data Warehouse Method”, Prentice Hall, 1999, ISBN:0-13-0812306-0.

3 David Heckerman, “Bayesian Networks for Data Mining”, Data Mining and Knowledge Discovery 1, 79–119 (1997)

4 `http://www.statserv.com`

5 P. Cheeseman, J. Kelly, M. Self, J. Stutz, W. Taylor and D. Freeman, “AutoClass: A Bayesian Classification System”, Proceedings of the Fifth International Conference on Machine Learning, Ann Arbor, MI. June 12-14 1988. Morgan Kaufmann Publishers, San Francisco, 1988, pp. 54-64

6 J. Goebel, K. Volk, H. Walker, F. Gerbault, P. Cheeseman, M. Self, J. Stutz and W. Taylor, “A Bayesian Classification of the IRAS LRS Atlas”, Astron. Astrophys. 222, L5-L8 (1989).

7 Michael J. Carey, David J. DeWitt, Chander Kant, Jeffrey F. Naughton, “A Status Report on the 007 OODBMS Benchmarking Effort”, OOPSLA 1994: 414-426

8 Michael J. Carey, David J. DeWitt, Jeffrey F. Naughton, “The 007 Benchmark”, SIGMOD Conference 1993: 12-21