AIM

Association for Information and Image Management

1100 Wayne Avenue, Suite 1100
Silver Spring, Maryland 20910
301/587-8202

Centimeter

Inches

MANUFACTURED TO AIM STANDARDS
BY APPLIED IMAGE, INC.

1 of 1

Database Computing for the SSC

Grant Number ER25133

Progress Report

Robert Grossman

June, 1993

1 Introduction

This is a progress report for the Grant "The PASS Project: Database Computing for the SSC" for the period September 1, 1993 – May 30, 1993. The number of the grant is ER25133; the institution is the University of Illinois at Chicago; and the principal investigator is Robert Grossman.

The PASS Project is broadly concerned with developing algorithms and software tools for making numerically intensive queries on very large amounts of data, especially high energy physics data. The approach is to view the analysis of such data as an appropriate query on a persistent object store. The PASS Project is a joint effort between Argonne National Laboratory, Lawrence Berkeley Laboratory, the Superconductor Super Collider Laboratory, the University of Maryland, and the University of Illinois at Chicago.

2 Report of Technical Progress

This section describes progress made during the report period.

Project 93-1. PTool32. In prior work, we developed a persistent object manager, now called PTool32, which we used as a basis for developing path planning algorithms for persistent stores of trajectory segments from nonlinear systems [3]. During the report period, we worked with other members of the project to develop an object data model for collider data, and the appropriate access methods so that a physicist could analyze collider data by making the appropriate query on a persistent store of collider events using PTool32. This work is described in [1]. As this report shows, software tools such as PTool32 provide significantly faster access to collider data than can be obtained using either current production technology or traditional relational database technology.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

MASTER

1
DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Project 93-2. PTool64/VTool. PTool32 uses memory mapping techniques to create, store and access persistent objects. This is the same technology that forms the basis for virtual memory. Currently, this technology is limited by the fact that most cpu's today can only reference objects within a 32 bit address space. Of course this will change, as 64 bit processors and operating systems become more widely available. For this reason, we modified PTool32 to use 64 bit id's for persistent objects. We also developed a companion software tool called VTool, which manages physical collections of persistent objects, called *folios*. When an object request is sent to PTool64, PTool64 first determines whether the object is currently available in persistent memory. If not, a fault is generated and the folio containing the object is located, either on a local cache disk, or by requesting a bitfile containing the folio from a hierarchical storage system, if necessary. Once the folio containing the object is available in local disk, a smaller unit called a *segment* containing the object is mapped to one of several slots in virtual memory managed by PTool64.

This approach combines the speed gain up provided by using memory techniques in a way which scales as the amount of data increases. It is also one of the first architectures which interfaces hierarchical storage systems and stores of which we are aware. We are currently working with other members of the project to access a persistent store of collider data using these tools. Preliminary results are available and look promising. This work is described in [4].

This work was also supported by NASA. With this support, stores of trajectory segments were created as a basis for path planning problems, and stores of vortex rings were created in order to study their stability numerically and visually.

Project 93-3. Reference Model. Working with other members of the project, we developed a preliminary version of a software reference model for this project. The model uses three primary concepts:

- objects
- processes
- nodes

Objects are an abstraction of data; processes are an abstraction of transformations of data; and nodes are an abstraction of a network supporting the storage and computation of data. These concepts provide enough of a framework so that not only can we model and study the problem by developing an appropriate data model, but we can also model and study the placement of data and its storage, accessing, and querying. This is a crucial step in our development of parallel and distributed query algorithms. This object model is currently used by us to model the collider data when creating and querying persistent stores of event data.

The object component of this software model is essentially OMG compliant. The other components are not covered by the OMG model. Our intention is to make the model ODMG compliant, when that model is released later this summer.

A report describing this model is nearing completion. This work was also supported by NSF and NASA.

Other Projects. During the report period, we also developed a number of benchmarks appropriate for our work and began to develop algorithms for making distributed and parallel queries on object stores. Reports describing this are currently in preparation.

3 Description of Planned Work for the Coming Year

This section describes our proposed work for the coming year. Further details are in the original proposal.

Project 94-1. PTool64/VTool. An important part of the proposed work for the coming year will focus on populating and accessing persistent object stores containing high energy physics event data—EventStores. We plan to extend the work described in the section above and to develop

- caching and migration algorithms;
- benchmarks;
- strategies for executing parallel and distributed queries, which are numerically intensive.

We plan to instrument PTool64 and VTool in order to obtain reliable profile information about the application specific benchmark queries that we have developed. With this information, we intend to develop adaptive caching and migration algorithms and distributed and parallel query algorithms suitable for persistent stores which access data from hierarchical storage systems.

On the basis of this work, we plan to redesign PTool64/VTool incorporating the lessons we learned. We plan to make PTool64/VTool CORBA and ODMG compliant beginning with this release.

Most of the work in this project will focus on developing appropriate query algorithms for accessing the EventStore. The algorithms will be encapsulated into appropriate classes and should work, with minimum modification, on any ODMG compliant object store. This has been our experience to date.

With support from NSF and NASA, we also plan to continue our work using stores of CFD data and stores of trajectory segments from nonlinear systems. It has been our experience that these stores exhibit different access patterns. It is our intention to work with a small number of application stores in order to increase our understanding of different object i/o behavior exhibited by scientific

computing. Our layered and modular tool based approach makes it relatively simple to work with a variety of different access algorithms depending upon the store.

Project 94-2. Interfacing stores to storage systems. Developing efficient and scalable interfaces from a persistent object store to a hierarchical storage system is an important part of our proposed software development. The hierarchical storage system necessarily includes distributed memories, disks, RAID's, and tertiary storage devices. We propose to continue our work designing a storage system-persistent object store interface that is compliant with the IEEE Storage System Reference Model.

We plan to continue our work with VTool, described in the section above, which is one strategy for interfacing stores and storage systems. VTool caches and migrates collections of objects (folios) to and from the storage system, but not entire stores. We have also gained some experience in caching and migrating entire stores to and from the storage system. One goal during this year will be to understand the trade-offs between these two strategies and whether some hybrid strategy can be designed.

Another important goal is to understand how using distributed file systems such as AFS affects the strategies for interfacing stores and storage systems.

We plan to coordinate this work with current file system research and with advances in network protocols for layered systems through our participation in the proposed Scalable Input Output Initiative.

Other Projects. We plan to continue other work at a lower level of effort:

- using the reference model to study analytically and via simulation different strategies for distributing and parallelizing queries;
- techniques for working with large numbers of collections and stores using metadata;
- and the concept of registering a collection so that it is available to other users by both its properties and its name.

4 Additional Information

4.1 Publications

Publications written during this report period appear at the end of the section.

4.2 Technology Transfer

During the past year, we have held three seminars explaining how companies can use object store technology as a component of their information systems. Approximately twenty companies participated in these seminars.

As a result of these seminars we have evolved the following technology transfer plan for the coming year. By making use of PTool/VTool, or one of the public domain persistent object managers available, interested companies can develop application specific tools to meet their own needs. As a first step in this direction, we have selected two companies who have shown interest in this approach: the plan is for each company to devote a software engineer to the project for a period of six months, with the expectation that this will be renewed for another six months. During this period, the software engineer will work with us on a general project task using the PTool/VTool technology. As a result of this experience, it should be relatively easy for the software engineer after returning to the company to undertake a company specific task using object manager technology. We expect to modify this technology transfer plan, depending upon our experience during the first year. We expect the first two software engineers to work with us in this fashion beginning sometime during the summer of 1993.

4.3 Cost Sharing

During the report period, we were awarded a three year NSF Grant Number IRI 9224605 to support the work of Robert Grossman and three co-principal investigators and their students to use PTool/VTool technology to explore issues in scientific visualization and incremental query processing. The amount of the award is approximately 135K per year.

During the report period, we received a continuation of our NASA award Number NAG2-513. This will be used to continue our research exploring the analysis of nonlinear control systems using database computing. The amount of the award is approximately 30K per year.

4.4 Collaborations

During the report period, we have:

1. collaborated with Mike Folk at NCSA to link PTool with HDF formatted data;
2. collaborated with Nancy Yeager at NCSA to link VTool with Unitree;
3. collaborated with John Guckenheimer at Cornell University to link PTool with dstool, a software tool for studying nonlinear systems developed at Cornell.

Several others have also made some preliminary evaluation of PTool/VTool for their work.

References

- [1] A. Baden, L. Cormell, C. T. Day, R. Grossman, P. Leibold, D. Lifka, D. Liu, S. Loken, E. Lusk, J. F. MacFarlane, E. May, U. Nixdorf, L. E. Price, X. Qin, B. Scipioni, and T. Song, "Database Computing in HEP—Progress Report," *Computing in High Energy Physics 1992*, to appear.
- [2] T. Defanti, R. Grossman, J. Leigh, D. Likfa, M. Papka, and X. Qin, "Analyzing high energy physics data with database computing," Showcase Presentation, 1992 SIGGRAPH, Chicago, July 29, 1992.
- [3] R. L. Grossman, S. Mehta, X. Qin, "Path planning by querying persistent stores of trajectory segments," *Laboratory for Advanced Computing Technical Report Number 93-R3*, University of Illinois at Chicago, 1993, to appear.
- [4] R. L. Grossman, D. Likfa, and X. Qin, "An object manager utilizing hierarchical storage," *Twelfth IEEE Symposium on Mass Storage Systems*, IEEE Press, Los Alamites, 1993, to appear.
- [5] E. May, D. Lifka, E. Lusk, L. E. Price, C. T. Day, S. Loken, J. F. MacFarlane, A. Baden, R. Grossman, X. Qin, L. Cormell, A. Gauthier, P. Leibold, J. Marstaller, U. Nixdorf, B. Scipioni, "Requirements for a Sytem to Analyze HEP Events Using Database Computing" *Twelfth IEEE Symposium on Mass Storage Systems*, IEEE Press, Los Alamites, 1993, to appear.
- [6] R. Grossman, X. Qin, C. T. Day, S. Loken, J. F. MacFarlane, E. May, D. Lifka, E. Lusk, L. E. Price, A. Baden, L. Cormell, P. Leibold, D. Liu, U. Nixdorf, B. Scipioni, and T. Song, "Database Challenges: Analyzing High Energy Physics Data Using Database Computing," *Laboratory for Advanced Computing Technical Report*, Number LAC 93-R7, University of Illinois at Chicago, December, 1993, to appear.

**DATE
FILMED**

8 / 09 / 93

END