

An Integrated Graphical User Interface for High Performance Distributed Computing

Xiaohui Shen, Wei-keng Liao and Alok Choudhary
Center for Parallel and Distributed Computing
Department of Electrical and Computer Engineering
Northwestern University
Evanston, IL 60208
{xhshen,wklio,choudhar}@ece.nwu.edu

Abstract

It is very common that modern large-scale scientific applications employ multiple compute and storage resources in a heterogeneously distributed environment. Working effectively and efficiently in such an environment is one of major concerns for designing meta-data management systems. In this paper, we present an integrated graphical user interface (GUI) that makes the entire environment virtually an easy-to-use control platform for managing complex programs and their large datasets. To hide the I/O latency when the user carries out interactive visualization, aggressive prefetching and caching techniques are employed in our GUI. The performance numbers show that the design of our Java GUI has achieved the goals of both high performance and ease-of-use.

1 Introduction

Modern simulations not only generate huge amounts of data, they also employ multiple techniques to process data generated: these processes include data analysis, visualization and so on. As data size for these applications is huge, large hierarchical storage system is used as data repository. In addition, databases are also introduced to make data management easily. Therefore, there are multiple resources involved in a modern large-scale scientific environment and these resources are heterogeneous and distributed. Without designing an efficient integrated environment, users have to deal with these resources manually and explicitly. In general, the requirements for the systems to support modern simulations are characterized as follows.

- **High performance simulation** For data intensive applications, state-of-the-art I/O optimizations such as collective I/O, prefetch, prestage and so forth should be employed to alleviate I/O problems of simulations.

- **High performance post-processings** If the user only considers simulation alone, it may result in bad performance when she carries out post-processings such as visualization and data analysis etc because these post-processings may not share the same access pattern with the simulations. The user should be careful to arrange the layouts of her datasets properly on storage systems.
- **Easy-to-use** Databases are employed to fulfill this purpose. The database can maintain meta-data information about the applications, datasets, storage systems and so on and it also provides easy query capabilities.
- **Integrated graphical environment** If the user works on an uniform platform rather than deal with distributed resources manually and explicitly, high efficiency can be achieved. Java proves to be a powerful language for such a task.
- **Latency reducing in the integrated environment** Given the fact that the speed of networks does not meet the user's requirements, aggressive optimizations such as prefetching and caching in a distributed environment are required to hide the network latency and reduce the probability of network failures.

A computational scientist would be overwhelmed in her application development if she has to consider by herself all these issues which are beyond her expertise. A lot of work addressing the above issues has been done separately in literature, few of them have considered them in a completely unified framework. Brown et al. [3] build a meta-data system on top of HPSS using DB2 from IBM. The SRB [2] and MCAT [12] provide a uniform interface for connecting to heterogeneous data resources over a network. Three I/O-intensive applications from the Scalable I/O Initiative Application Suite are studied in [8]. But all these works only address one aspect of the issues we discussed above. Our previous work [6] is a first step toward considering multiple factors in a complete picture. We have designed an active meta-data management system that takes advantage

of state-of-the-art I/O optimizations as well as maintaining ease-of-use features by employing relational database techniques. In this paper, we present further considerations about integrated environment and its optimizations in a distributed environment based on our previous work. We make the following contributions:

- Present a high performance data management system in which database performs an active role in making I/O strategies, as well as managing huge amounts of data easily.
- Present the design of an integrated graphical environment in Java for large-scale scientific applications in a distributed environment. In our unified framework, users work only on their local machines and our GUI hides all the details of distributed resources. Users can launch the parallel application, carry out data analysis and visualization, query databases and browse the tables in an uniform interface.
- Present an I/O latency reducing scheme which significantly improves I/O response time when the user carries interactive visualization.

The remainder of the paper is organized as follows. In Section 2 we introduce an astrophysics application and a parallel volume rendering application that we used in our work. A shorthand notation is also introduced for convenience. In Section 3 we give an overview of our design of meta-data systems (MDMS). We present our design of the integrated Java graphical environment for scientific simulations in Section 4. The functions and the inner-mechanisms that GUI provides are presented. In Section 5 we present our I/O latency reducing techniques. Finally we conclude our paper in Section 6.

2 Introduction to Applications

Our first application, called Astro-3D or astro3d [1] henceforth, is a code for scalably parallel architectures to solve the equations of compressible hydrodynamics for a gas in which the thermal conductivity changes as a function of temperature ¹.

From computer system's point of view, the application just generates a sequence of data and dumps them on storage. Later, the user may visualize the datasets in which she

¹The code has been developed to study the highly turbulent convective envelopes of stars like the sun, but simple modifications make it suitable for a much wider class of problems in astrophysical fluid dynamics. The algorithm consists of two components: (a) a finite difference higher-order Godunov method for compressible hydrodynamics, and (b) a Crank-Nicholson method based on nonlinear multigrid method to treat the nonlinear thermal diffusion operator. These are combined together using a time splitting formulation to solve the full set of equations

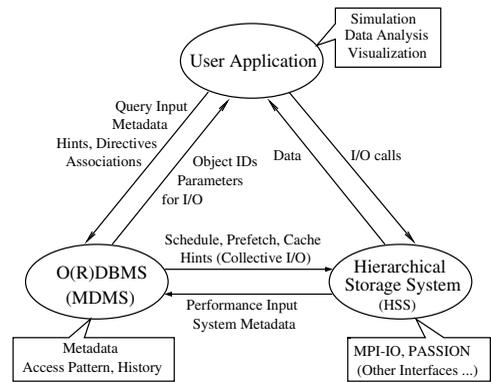


Figure 1. Three-tiered architecture.

may be interested. In Astro-3D for example, it generates six datasets such as temperature, pressure, etc. for each iteration in one run. Each of these datasets is written in a single file. Therefore, a data file is uniquely identified by dataset name, run id and the iteration number. We make the following notations to express a data file by concatenating the dataset name, run id and iteration number: *dataset-runid-iteration*. For example, if the temperature is dumped at the first iteration in the fifth run, it is notated as *temperature-5-1*; if the pressure is dumped at the second iteration in the sixth run, it is expressed as *pressure-6-2* and so on.

Our second application is a parallel volume rendering code (called volren henceforth). It generates a 2D image by projection given a 3D input file. The algorithm is described as follows: The data points of the input file are at the corners of cubic cells. The entire data volume is a rectangular solid made up of these cells. An orthographic projection is done through the volume parallel with one of its three axes. One of the 2 faces of the data volume that are vertical to the projection is used as the image. If there are $M \times N$ data points in this face, there are $(M-1) \times (N-1)$ pixels in the image. There is one line of projection per pixel. The densities in the stack of cells along this line are added (until their sum is 1.0), and the result becomes the alpha and grey value for that pixel. The command-line arguments include number of iterations per run, sizes of input files and so on.

From computer system's point of view, again, volren just reads a 3D input file and creates an 2D image file for each iteration per run. For example, volren may generate four image files (*image-5-1*, *image-5-2*, *image-5-3*, *image-5-4*) given four input data files at the fifth run.

3 Design of Meta-data Management System (MDMS)

Figure 1 shows a novel architecture we proposed in [6]. The three-tiered architecture contains three key com-

ponents: (1) parallel application, (2) meta-data management system (MDMS), and (3) hierarchical storage system (HSS). These three components can co-exist in the same site or can be fully-distributed across distant sites. The MDMS is an active part of the system: it is built around an OR-DBMS [16, 15] and it mediates between the user program and the HSS. The user program can send query requests to MDMS to obtain information about data structures that will be accessed. Then, the user can use this information in accessing the HSS in an optimal manner, taking advantage of powerful I/O optimizations like collective I/O [17, 5, 11], prefetching [10], prestaging [7], and so on. The user program can also send *access pattern hints* to the MDMS and let the MDMS to decide the best I/O strategy considering the storage layout of the data in question. These access pattern hints span a wide spectrum that contains inter-processors I/O access patterns, information about whether the access type is read-only, write-only, or read/write, information about the size (in bytes) of average I/O requests, and so on. In this section, we give an overview of our MDMS designs.

Our MDMS tries to keep meta-data information summarized as follows.

- It stores information about the abstract storage devices (ASDs) that can be accessed by applications. By querying the MDMS² the applications can learn where in the HSS their data reside (i.e., in what part of the storage hierarchy) without the need of specifying file names. They can also access the performance characteristics (e.g., speed) of the ASDs and select a suitable ASD (e.g., a disk subsystem consisting of eight separate disk arrays) to store their datasets.

- It stores information about the *storage patterns* and *access patterns* of data sets. For example, a specific multi-dimensional array that is striped across four disk devices in round-robin manner will have an entry in the MDMS. The MDMS utilizes this information in a number of ways. The most important usage of this information, however, is to decide a parallel I/O method based on *access patterns (hints)* provided by the application. By *comparing* the storage pattern and access pattern of a dataset, the MDMS can, for example, advise the HSS to perform collective I/O [9] or prefetching [10] for this dataset.

- It keeps meta-data for specifying access history and trail of navigation (not covered in this abstract).

Notice that the MDMS is not merely a data repository but also an *active* component in the overall data management process. It *communicates* with applications as well as the HSS and can *influence* the decisions taken by the both.

Figure 2 shows the MDMS database tables and Figure 3

²These queries are performed using user-friendly constructs. It would be very demanding to expect the user to know SQL [14] or any other query language.

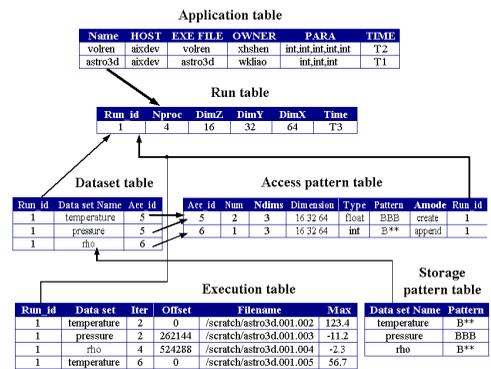


Figure 2. Internal representation in MDMS.

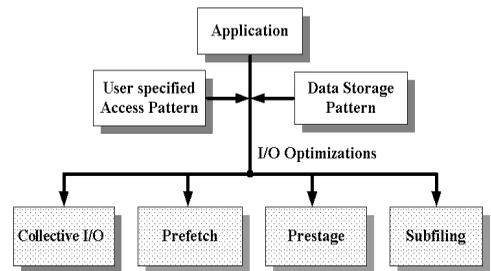


Figure 3. Selecting an I/O optimization.

shows the MDMS internal optimization mechanisms.

4 Design of Java Graphical User Interface

4.1 Architecture of Integrated Java GUI

As it is distributed in nature, our programming environment involves multiple resources across distant sites. Consider our current working environment, we are working on local HP or SUN workstations, the visualization tools are installed on a Linux machine, our database (POST-GRESQL) is located on another machine and our parallel applications run on a 16 node IBM SP2. Although these machines are within our department, they could be distributed across any places on Internet.

When a user starts to work in such a distributed environment, she needs to go through the following procedures:

- (1) log on to SP2 and submit the parallel application.
- (2) When the application is finished, she needs to log on to the database host and use native SQL language to inspect the database to find datasets she would be interested in for visualization.
- (3) When the user is interested in a particular dataset, she would transfer the data file explicitly, for example us-

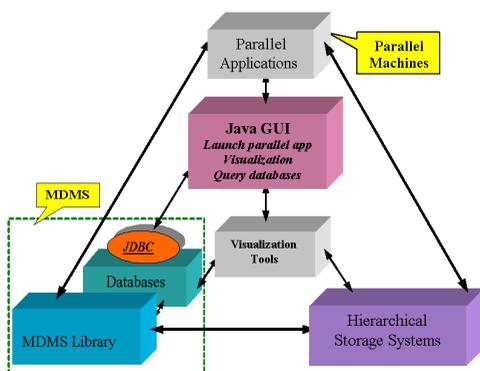


Figure 4. Java GUI in overall system.

ing ftp, from SP2 where data are located to the visualization host where visualization tools reside.

- (4) Log on to the visualization host (DATA) and start the visualization process.
- (5) Repeat steps 2-4 as long as there exist datasets to be visualized.

Obviously, these steps might be very time-consuming and inconvenient for the users. To overcome this problem (which is due to the distributed nature of the environment), an integrated Java graphical user interface (GUI) is implemented and integrated to our application development environment. The goal of the GUI is to provide users with an integrated graphical environment that hides all the details of interaction among multiple distributed resources (including storage hierarchies).

We use Java because Java itself proves to be a key enabling access language and operating environment with support for all of our platforms of interest, including IBM AIX, Linux, Windows NT, Solaris, and others. Transparency is made possible by the many platform independent abstractions of Java, including process management (a built-in class), multithreading (a language feature), networking and streams (built-in classes), GUI components (the Abstract Windowing Toolkit), and database access (JDBC). In this environment, the users need to work only with GUI locally, rather than go to different sites to submit parallel applications or to do file transfers explicitly. Figure 4 shows how GUI is related to other parts of our system. It actively interacts with three major parts of our system:

- Interacts with parallel machines to launch parallel applications.
- Interacts with databases through JDBC to help users query meta-data from databases.
- Interacts with visualization tools to carry out visualization process.

4.2 Main functions of GUI

The main functions that GUI provides are described as follows:

- **Registering new applications** To start a new application, the user needs to create a new suite of tables for the new application. By GUI, the user needs only to specify attributes of run table that she would be interested in, and all the other tables will be created automatically with run table.
- **Running applications remotely** The applications are usually running somewhere on parallel machines such as SP2, which are specified by the user when she registers a new application. Therefore, remote shell command is used in GUI to launch the job on remote parallel machines. The user can also specify command line arguments in the small text fields. Defaults are provided and the user can change them as needed. The running results will be returned in the large text area. application: four processors are used and the sizes of datasets are 16*16*16.
- **Data Analysis and Visualization** Users can also carry out data analysis and visualization through our GUI. Data Analysis may come in a variety of flavors, it is quite application specific. For some applications, data analysis may simply calculate the maximum, minimum or average value of a given dataset, for some others, it may be plugged into the application and calculate the difference between two datasets and decide whether the dataset should be dumped now or later. Our system's current method of data analysis is to calculate the maximum, minimum and means of each dataset generated. From the GUIs point of view, it is no different than just submitting a remote job. Visualization is becoming an important approach in large-scale scientific simulation to inspect the inside nature of datasets. It is often a little more complicated than data analysis: first of all, the users' interest in a particular data set may be very arbitrary. Our approach is to list all the candidate datasets by searching the database by user-specified characteristics such as maximum, minimum, means, iteration numbers, pattern, mode and so on. Then the candidates are presented in radio box for user to choose easily. Second, the datasets that are created by parallel machines, are located either at parallel machines or stored in hierarchical storage systems. But our visualization tools are installed at other places. Therefore, inside GUI, we transparently transfer the data from the remote parallel machine or hierarchical storage systems to the visualization host's local disks and then start the visualization process. This is implemented by having a server

program written in C running on the storage side and the GUI (in Java) serves as the client. The user does not need to check the database for interesting datasets or do data transfer explicitly. Our current visualization tools include Visualization Toolkit (VTK), Java 3D, XV etc.

- **Table browsing and searching** Advanced users may want to search the database to find datasets of particular interest. So the table browsing and searching functions are provided in our GUI. The user can just move the mouse and pick a table to browse and search the data without logging on to a database host and typing native SQL script.

5 I/O Latency Reducing for Interactive Visualization

Our GUI can help the user work at a single site without consulting a variety of distributed components explicitly. A potential problem in this environment, however, is the I/O response time when the user carries out interactive visualization (VTK) on a sequence of datasets. Note that visualization tools such as VTK etc are usually installed on sequential workstations, while the datasets generated by parallel applications may be stored on the disks of parallel computing systems or on a remote storage system such as tertiary storage systems like HPSS [7]. Therefore, these datasets need to be first transferred to the local site from remote storage. As the data size could be large and the physical distance could be long, the I/O response time is significant. It is quite annoying that the user has to wait for tens of seconds for the availability of the data after she has launched the visualization process on a dataset. To reduce the I/O latency, a naive approach is that the user explicitly transfers all the data to the local disks first and then starts visualization process. The problems of this approach, however, are three-fold. First, the user has to deal with specific file names and locations manually which are very inconvenient; second, the user may not be sure what datasets she might use in the future and finally, the overall time is still long since the I/O time is not overlapped with the visualization process.

In this section, we present a scheme to address the problems of naive approach. By using prefetching and caching techniques, our approach can effectively hide the I/O latency, improving overall visualization performance in GUI. The basic idea of our approach is that when the user is working on a dataset (visualization), we make a prediction on what the next dataset that might be accessed, then spawn another thread to perform remote I/O and caches it on the local disks. Since visualization process with human interaction is slow, there is ample time to overlap prefetching.

In addition, the I/O response time for the next dataset can be significantly reduced when the required dataset can be serviced on local disk cache.

We use database to keep track of data locations and access history. Two tables are created in the database. One table is called *data-access-trace*, whose attributes (fields) include application name, dataset name, iteration number, date and time, run id etc. This table keeps all the datasets visualized by the user. Another table is called *data-cache*, whose attributes (fields) include, application name, dataset name, iteration number, local directory and reference counter. It keeps information about what datasets are currently cached on the local disks and how many times they are accessed by the user.

The key issue for prefetching is to decide which dataset should be fetched. In Section 2 we introduced an astrophysics application which is a representative of many scientific applications. We can see that many scientific applications generate time-serial datasets (each dataset is stored in a separate file) and these datasets will also be accessed by visualization tools one by one in time order. For example, by visualizing datasets ‘temperature’ from time step 0 to 20, the user can know how temperature changes as simulation goes on. This characteristic is good for prediction, because the next time step dataset is the best candidate for prefetching. Figure 5 (column 1,2) shows the results of prefetching. The I/O response time can be reduced significantly.

As some applications may take many time steps and the user may not always visualize all of them. For example, to have a quick view of how dataset ‘temperature’ changes as simulation goes on, the user may only pick part of data files in a strided manner such as temperature-9-0, temperature-9-2, temperature-9-4, etc. In this case, the naive prefetching would not help. To address this problem, we proposed an ‘adaptive stride’ scheme. In this approach, the previous stride is used to predict next dataset. For example, after datasets temperature-9-0 and temperature-9-2 are accessed, we predict the next dataset should be temperature-9-4 since the last stride is 2. Figure 5 (column 3,4,5) shows the results. The adaptive approach dramatically reduces I/O latency. For some cases, the user may be interested in multiple datasets. For example, the changes of temperature and pressure may influence with each other, so the user may inspect temperature and pressure datasets alternatively. The adaptive stride prefetching can still deliver significant performance improvement in Figure 5 (column 6,7,8).

Another scenario could happen is that the datasets to be visualized may not have fixed stride. By checking the data access history in table *data-access-trace* of previous runs, we can still make a good prediction since the user may change run-time parameters for each run and may still interested in the same set of datasets. Figure 5 (column 9, 10) shows the results.

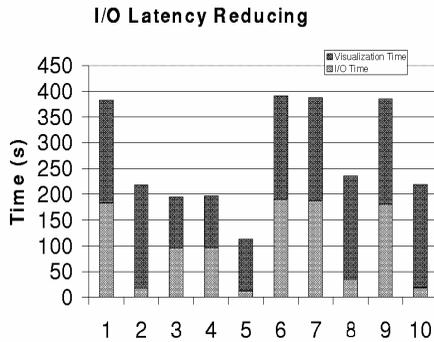


Figure 5. I/O Latency Reducing by Prefetching for Astrophysics data visualization. The data is located at disks of SP2 and visualization tool (VTK) is installed on a Linux machine. The size of each dataset is 8MB. The average visualization time on a dataset is about 10s. (1) No prefetch, Input A; (2) Prefetch, Input A; (3) No prefetch, Input B; (4) Prefetch (fixed stride), Input B; (5) Prefetch (adaptive stride), Input B; (6) No prefetch, Input C; (7) Prefetch (fixed stride), Input C; (8) Prefetch (adaptive stride), Input C; (9) No prefetch, Input D; (10) Prefetch (according to previous runs), Input E. Input A = (temp-9-0, temp-9-1, temp-9-2, ..., temp-9-20), B = (temp-9-0, temp-9-2, temp-9-4, ..., temp-9-20), C = (temp-9-0, press-9-0, temp-9-2, press-9-2, ..., temp-9-20, press-9-20); D = (temp-9-1, temp-9-5, temp-9-4, temp-9-5, temp-9-10, ...), E = (temp-10-1, temp-10-5, temp-10-4, temp-10-5, temp-10-10, ...).

6 Conclusions

In this paper, we presented an integrated Java graphical user interface (GUI) to efficiently help users work on an environment that is characterized by distributed and heterogeneous natures. Our GUI provides users an unified interface to all the resources and platforms presented to large-scale scientific applications. In addition, an I/O response time reducing technique has been integrated into our GUI to hide I/O latency for interactive visualization. All these works take advantage of Java's powerful features such as platform independence, portability, process management, multithreading, networking and streams. The database is also plays an important role which makes the whole framework complete. In the future, we would investigate other optimizations in our environment, such as subfiling [13]. The relationship between prefetch and caching [4] in our context will also be studied.

References

- [1] A. Malagoli, A. Dubey, and F. Cattaneo. A Portable and Efficient Parallel Code for Astrophysical Fluid Dynamics. <http://astro.uchicago.edu/Computing/On-Line/cfd95/camelse.html>
- [2] C. Baru, R. Moore, A. Rajasekar, and M. Wan. The SDSC storage resource broker. In *Proc. CASCON'98 Conference*, Dec 1998, Toronto, Canada.
- [3] P. Brown, R. Troy, D. Fisher, S. Louis, J. R. McGraw, and R. Musick. Meta-data sharing for balanced performance. In *Proc. the First IEEE Meta-data Conference*, Silver Spring, Maryland, 1996.
- [4] P. Cao, E. Felten, and K. Li. Application-controlled file caching policies. In *Proc. the 1994 Summer USENIX Technical Conference*, pages 171–182, June 1994.
- [5] A. Choudhary, R. Bordawekar, M. Harry, R. Krishnaiyer, R. Ponnusamy, T. Singh, and R. Thakur. PASSION: parallel and scalable software for input-output. *NPAC Technical Report SCCS-636*, Sept 1994.
- [6] A. Choudhary, M. Kandemir, H. Nagesh, J. No, X. Shen, V. Taylor, S. More, and R. Thakur. Data management for large-scale scientific computations in high performance distributed systems, In *Proc. the 8th IEEE International Symposium on High Performance Distributed Computing (HPDC'99)*, August 3-6, 1999, Redondo Beach, California.
- [7] R. A. Coyne, H. Hulén, and R. Watson. The high performance storage system. In *Proc. Supercomputing 93*, Portland, OR, November 1993.
- [8] P. E. Crandall, R. A. Ayt, A. A. Chien, and D. A. Reed. Input/output characteristics of scalable parallel applications. In *Proceedings of Supercomputing '95*.
- [9] J. del Rosario, R. Bordawekar, and A. Choudhary. Improved parallel I/O via a two-phase run-time access strategy. In *Proc. the 1993 IPPS Workshop on Input/Output in Parallel Computer Systems*, April 1993.
- [10] C. S. Ellis and D. Kotz. Prefetching in file systems for MIMD multiprocessors. In *Proc. the 1989 International Conference on Parallel Processing*, pages 1:306–314, St. Charles, IL, August 1989. Pennsylvania State Univ. Press.
- [11] D. Kotz. Disk-directed I/O for MIMD multiprocessors. In *Proc. the 1994 Symposium on Operating Systems Design and Implementation*, pages 61–74. USENIX Association, Nov 1994.
- [12] MCAT <http://www.npaci.edu/DICE/SRB/mcat.html>.
- [13] G. Memik, M. Kandemir, A. Choudhary, Valerie E. Taylor. APRIL: A Run-Time Library for Tape Resident Data. To appear in *the 17th IEEE Symposium on Mass Storage Systems*, March 2000.
- [14] R. Ramakrishnan. *Database Management Systems*, The McGraw-Hill Companies, Inc., 1998.
- [15] M. Stonebraker. *Object-Relational DBMSs : Tracking the Next Great Wave*. Morgan Kaufman Publishers, ISBN: 1558604529, 1998.
- [16] M. Stonebraker and L. A. Rowe. The design of Postgres. In *Proc. the ACM SIGMOD '86 International Conference on Management of Data*, Washington, DC, USA, May 1986, pp. 340–355.
- [17] R. Thakur, W. Groppe, and E. Lusk. Data sieving and collective I/O in ROMIO. To appear in *Proc. the 7th Symposium on the Frontiers of Massively Parallel Computation*, February 1999.
- [18] *UniTree User Guide*. Release 2.0, UniTree Software, Inc., 1998.