

# A High-Performance Application Data Environment for Large-Scale Scientific Computations

Xiaohui Shen, Wei-keng Liao, Alok Choudhary, *Senior Member, IEEE*,  
Gokhan Memik, and Mahmut Kandemir, *Member, IEEE*

**Abstract**—Effective high-level data management is becoming an important issue with more and more scientific applications manipulating huge amounts of secondary-storage and tertiary-storage data using parallel processors. A major problem facing the current solutions to this data management problem is that these solutions either require a deep understanding of specific data storage architectures and file layouts to obtain the best performance (as in high-performance storage management systems and parallel file systems), or they sacrifice significant performance in exchange for ease-of-use and portability (as in traditional database management systems). In this paper, we discuss the design, implementation, and evaluation of a novel application development environment for scientific computations. This environment includes a number of components that make it easy for the programmers to code and run their applications without much programming effort and, at the same time, to harness the available computational and storage power on parallel architectures.

**Index Terms**—Data intensive computing, access pattern, storage pattern, MDMS.

## 1 INTRODUCTION

EFFECTIVE data management is a crucial part of the design of large-scale scientific applications. An important subproblem in this domain is to optimize the data flow between parallel processors and several types of storage devices residing in a storage hierarchy. While a knowledgeable user can manage this data flow by exerting a great effort, this process is time-consuming, error-prone, and not portable.

To illustrate the complexity of this problem, we consider a typical computational science analysis cycle shown in Fig. 1. As can be seen easily, in this cycle, there are several steps involved. These include mesh generation, domain decomposition, simulation, visualization and interpretation of results, archiving of data and results for postprocessing and check-pointing, and adjustment of parameters. Consequently, it may not be sufficient to consider simulation alone when determining how to store or access data sets, because these data sets are used in other steps as well. In addition, these steps may need to be performed in a heterogeneous distributed environment and the data sets in question can be persistent on secondary or tertiary storage.

Among the important issues in this analysis cycle are detection of I/O access patterns for data files, determination of suitable data storage patterns, and effective data analysis and visualization.

Obviously, designing effective I/O strategies in such an environment is not particularly suitable for a computational scientist. To address this issue, over the years, several solutions have been designed and implemented. While each of these solutions is quite successful for a class of applications, we feel that the growing demand for large-scale data management necessitates novel approaches that combine the best characteristics of current solutions in the market. For example, parallel file systems [11], [30], [9] might be effective for applications whose I/O access patterns fit a few specific forms. They achieve impressive performance for these applications by utilizing smart I/O optimization techniques such as prefetching [19], caching [24], [6], and parallel I/O [17], [12]. However, there are serious obstacles preventing the parallel file systems from becoming a global solution to the data management problem. First of all, user interfaces of the file systems are, in general, low-level [22], allowing the users to express access patterns of their applications using only low-level structures such as file pointers and byte offsets. Second of all, nearly every file system has its own suite of I/O commands, rendering the process of porting a program from one machine to another a very difficult task. Third, the file system policies and optimization parameters are, in general, hard-coded within the file system and, consequently, work for only a small set of access patterns. While runtime systems and libraries like MPI-IO [10], [35] and others [37], [3], [8] present users with higher-level, more structured interfaces, the excessive number of calls to select from, each with several parameters, make the user's job

- X. Shen is with the Core Technology Department, Motorola Inc., 1000 Technology Way, Libertyville, IL 60048. E-mail: axs095@email.mot.com.
- W.-k. Liao, A. Choudhary, and G. Memik are with the Electrical and Computer Engineering Department, Northwestern University, L359 Technological Institute, 2145 Sheridan Road, Evanston, IL 60208-3118. E-mail: {wkliao, choudhar, memik}@ece.northwestern.edu.
- M. Kandemir is with the Department of Computer Science and Engineering, The Pennsylvania State University, University Park, PA 16802. E-mail: kandemir@cse.psu.edu.

Manuscript received 18 Dec. 2001; revised 24 Nov. 2002; accepted 20 May 2003.

For information on obtaining reprints of this article, please send e-mail to: [tpds@computer.org](mailto:tpds@computer.org), and reference IEEECS Log Number 115578.




Fig. 1. A typical computational science analysis cycle.

very difficult. Also, the usability of these libraries depends largely on how well user's access patterns and library calls' functionality match [21].

An alternative to parallel file systems and runtime libraries is database management systems (DBMS). They present a high-level, easy-to-use interface to the user and are portable across a large number of systems including SMPs and clusters of workstations. In fact, with the advent of object-oriented and object-relational databases [32], they also have the capability of handling large data sets such as multidimensional arrays and image/video files [15]. A major obstacle in front of DBMS (as far as the effective high-level data management is concerned) is the lack of powerful I/O optimizations that can harness parallel I/O capabilities of current multiprocessor architectures. In addition to that, the data consistency and integrity semantics provided by almost all DBMS put an added obstacle to high-performance. Finally, although hierarchical storage management systems (e.g., [39]) are effective in large-scale data transfers between storage devices in different levels of a storage hierarchy, they also, like parallel file systems and DBMS, lack application specific access pattern information and, consequently, their I/O access strategies and optimizations are targeted at only a few well-defined access and storage patterns.

In this paper, we present a novel application development environment for large-scale scientific applications that manipulate secondary storage and tertiary storage resident data sets. Our primary objective is to combine the advantages of parallel file systems and DBMS without suffering from their disadvantages. To accomplish this objective, we designed and implemented a multicomponent system that is capable of applying state-of-the-art I/O optimizations without putting an excessive burden on users. Our contributions are summarized as follows:

- We present a *metadata management system*, called MDMS, that keeps track of I/O accesses and enables suitable I/O strategies and optimizations depending on the access pattern information. Unlike classical user-level and system-level metadata systems [18], [28], the main reason for the existence of MDMS is to keep performance-oriented metadata and utilize these metadata in deciding suitable I/O strategies.
- We explain how the MDMS interacts with parallel applications and hierarchical storage systems (HSS), relieving the users from the low-level management of data flow across multiple storage devices. In this

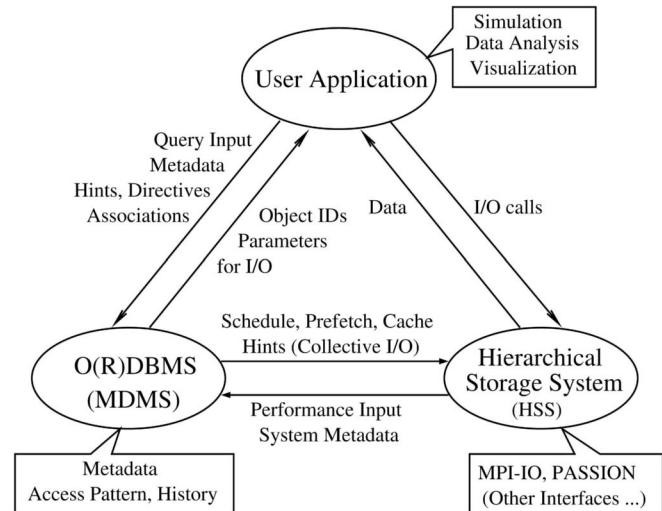


Fig. 2. Three-tiered architecture.

respect, the MDMS plays the role of an easy-to-use interface between applications and HSS.

- We present a tape device-oriented optimization technique, called *subfiling*, that enables fast access to small portions of tape-resident data sets and show how it fits in the overall application development environment.
- We illustrate how data analysis and visualization tools can be integrated in our environment.
- We propose an automatic code generator component (ACG) to help users utilize the metadata management system when they are developing new applications.
- We present an integrated Java graphical user interface (IJ-GUI) that makes the entire environment virtually an easy-to-use control platform for managing complex programs and their large data sets.
- We present performance numbers from our initial implementation using four I/O-intensive scientific applications.

The core part of our environment is a three-tiered architecture shown in Fig. 2. In this environment, there are three key components: 1) parallel application, 2) metadata management system (MDMS), and 3) hierarchical storage system (HSS). These three components can coexist 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 Object-Relational DBMS (OR-DBMS) [33], [32] and it mediates between the user program and the HSS. OR-DBMS is a marriage of the SQL from the relational world and the modeling primitives from the object world, and is very effective for large and complex scientific data. 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 [36], [8], [23], prefetching [19], prestaging [14], etc. 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 interprocessors 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, etc. We believe that this is one of the first studies evaluating the usefulness of passing a large number of user-specified hints to the underlying I/O software layers. In this paper, we focus on the design of MDMS, including the design of database schema and MDMS library (user interface), the optimizations for tape-resident data sets, and an integrated Java graphical user interface (IJ-GUI) to help users efficiently work in our distributed programming environment. Our environment is different from previous platforms (e.g., [25], [2], [1], [5]) in that it provides intelligent data access methods for disk and tape-resident data sets.

The remainder of the paper is organized as follows: In Section 2, we present the design details of metadata management system including design of database tables and high-level MDMS library (user API). In Section 3, an optimization method to access tape-resident data sets is presented. In Section 4, we present an integrated Java graphical user interface (IJ-GUI) to assist users in distributed environments. In Section 5, our initial performance results are presented. In Section 6, we review the previous work on I/O optimizations. Finally, we conclude the paper and briefly discuss ongoing and future work in Section 7.

## 2 DESIGN OF METADATA MANAGEMENT SYSTEM (MDMS)

The metadata management system is an active middle-ware built at Northwestern University with the aim of providing a uniform interface to data-intensive applications and hierarchical storage systems. Applications can communicate with the MDMS to exploit the high-performance I/O capabilities of the underlying parallel architecture. The main functions fulfilled by the MDMS can be summarized as follows:

- It stores information about the abstract storage devices (ASDs) that can be accessed by applications. By querying the MDMS,<sup>1</sup> the applications can learn where in the HSS their data sets reside (i.e., in what parts of the storage hierarchy) without the need of specifying file names. They can also access the performance characteristics (e.g., speed, capacity, bandwidth) of the ASDs and select a suitable ASD (e.g., a disk subsystem consisting of eight separate disk arrays or a robotic tape device) to store their data sets. Internal data structures used in the MDMS map ASDs to physical storage devices (PSDs) available in the storage hierarchy.
- It stores information about the *storage patterns* (storage layouts) of data sets. For example, a specific multidimensional array that is striped across four disk devices in round-robin manner will have an entry in the MDMS indicating its storage pattern. 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 data set, the MDMS can, for example, advise the HSS to perform collective I/O [16] or prefetching [19] for this data set.

- It stores information about the *pending* access patterns.<sup>2</sup> It utilizes this information in taking some global data movement decisions (e.g., file migration [39], [14] and prestaging [39], [14]), possibly involving data sets from multiple applications.
- It keeps metadata for specifying access history and trail of navigation. This information can then be utilized in selecting appropriate optimization policies in successive runs.

Overall, the MDMS keeps vital information about the data sets and the storage devices in the HSS. Note 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 both.

The MDMS design consists of the design of database tables and the design of a high-level MDMS API. The database tables keep the metadata that will be utilized in performance-oriented I/O optimizations. The MDMS API, on the other hand, presents an interface to the clients of the MDMS. They are described in the subsequent subsections.

### 2.1 MDMS Tables

We have decided that, to achieve effective I/O optimizations automatically, the MDMS should keep five (database) tables for each application. These are *run table*, *storage pattern table*, *access pattern table*, *data set table*, and *execution table*. Since, in our environment, a single user might have multiple applications running, sharing tables among different applications would not be a good implementation choice because it might slow down the query speed when tables become large. In our implementation, we construct a table name by concatenating the application name and a fixed, table-specific name. Consequently, each application has its own suite of tables. For example, in an astrophysics application (called *astro3d* henceforth), the table names are *astro3d-run-table*, *astro3d-access-pattern-table*, etc., while in a parallel volume rendering application (called *volren* henceforth), they are *volren-run-table*, *volren-access-pattern-table*, and so forth. The tables with same fixed table name (e.g., *data set table*) have the same attributes for different applications except the run table, which is application specific: The user needs to specify interesting attributes (fields) for a particular application in the run table. For example, in *astro3d*, the run table may contain the number of dimensions and the dimension sizes of each array, the total number of iterations, the frequency of dumping for data analysis, the frequency of check-point dumping, etc. The functionality of each table is briefly summarized in Table 1. Note that, among these tables, the execution table is the most frequently updated one. It is

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

2. The pending access patterns here refer to future access patterns that will be known at runtime where a user can specify the type of access pattern that will occur. Using this information, the system can perform optimizations such as prefetching, prestaging from tape, data migration, etc.

TABLE 1  
Functionality of Database Tables Maintained in the MDMS

Table Name	Functionality	Primary Key
Run table	Records each run of the application with user-specified attributes	Run id
Dataset table	Keeps information about the datasets used each run	run id + association id
Access pattern table	Keeps the access pattern specified by user for each dataset	run id + dataset name
Storage pattern table	Keeps information on how data stored for each dataset	dataset name
Execution table	Records I/O activities of the run, including file path and name, offset, etc.	run id + dataset + iteration number

typically updated whenever the application in question dumps data on disk/tape for visualization and data analysis purposes. The run table, on the other hand, is updated once for each run (assigning a new run-id to each run). The data set table keeps the relevant information about data sets in the application, the access pattern table maintains the access pattern information, and the storage pattern table keeps information about storage layouts of the data sets. An advantage of using an OR-DBMS [33] in building the MDMS is being able to use pointers that minimize metadata replication, thereby keeping the database tables in manageable sizes. The MDMS also has a number of global (interapplication) tables to manage all applications, such as *application table*, which records all the application names, their host machines, etc., in the system, *visualization table*, where location of visualization tools can be found, and *storage devices table*, which maps ASDs to PSDs. An example use of our five database tables is illustrated in Fig. 3. Note that the memory requirements for MDMS are small. A typical run of a large scale application results in about 20-30K increase of metadata.

## 2.2 MDMS API

The MDMS API, which consists of a number of MDMS functions, is in the center of our programming environment.

Through this API, the programs can interact with the database tables without getting involved with low-level SQL-like commands. Our MDMS library is built on top of MPI-I/O [10], the emerging parallel I/O standard. MPI-I/O provides many I/O optimization methods such as collective I/O, data sieving, asynchronous I/O, etc. But, for most computational scientists with little knowledge of I/O optimizations and storage devices, it is very hard to choose the appropriate I/O routines from among numerous complicated MPI-I/O functions. Our MDMS API helps users choose the most suitable I/O functions according to user-specified data access pattern information. In this environment, an access pattern for a data set is specified by indicating how the data set is to be shared and accessed by parallel processors. For example, an access pattern such as (Block, \*) says that the two-dimensional data set in question is divided (logically) into groups of rows and each group of rows will be accessed by a single processor. These patterns are also used as storage patterns. As an example, for a two-dimensional disk-resident array, a (Block, \*) storage pattern corresponds to row-major storage layout (as in C), a (\*, Block) storage pattern corresponds to column-major storage layout (as in Fortran), and a (Block, Block) storage pattern corresponds to blocked storage layout which might be very useful for large-scale linear algebra applications whose data sets are amenable to blocking [37]. Our experience with large-scale, I/O-intensive codes indicates that, usually, the users know how their data sets will be used by parallel processors; that is, they have sufficient information to specify suitable access patterns for the data sets in their applications. Note that conveying an access pattern to the MDMS can be quite useful, as the MDMS can compare this access pattern with the storage pattern of the data set (which is kept in the storage pattern table), and can decide an optimal I/O access strategy.<sup>3</sup> For instance, an example use of this information might occur in the following way. If the user is going to access a data set in a (Block, Block) fashion while the data set is stored, say in a file on disk, as (Block, \*), the MDMS will automatically choose the MPI-I/O collective I/O function to

3. Our MDMD supports many regular partitioning patterns such as (Block, \*), (Block,Block), (Block, CYCLIC), etc. In addition, our recent work would allow users to specify irregular and more flexible access patterns by employing MPI-I/Os derived datatypes mechanism [34]. In this work, however, only the most popular patterns such as (Block, \*), (Block, Block) have been used for our experiments.

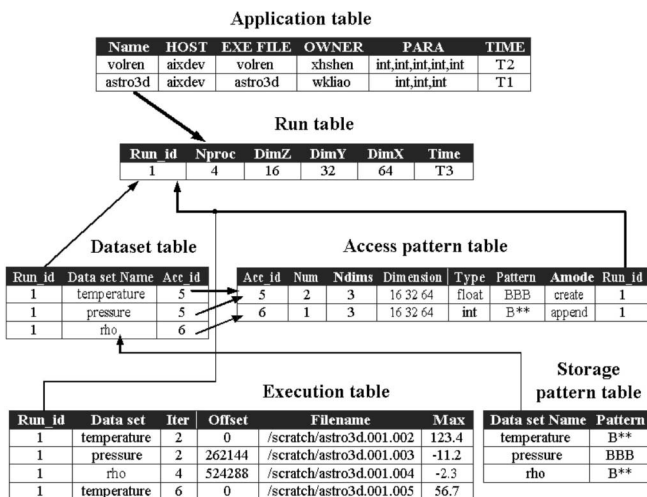


Fig. 3. Internal representation in MDMS.

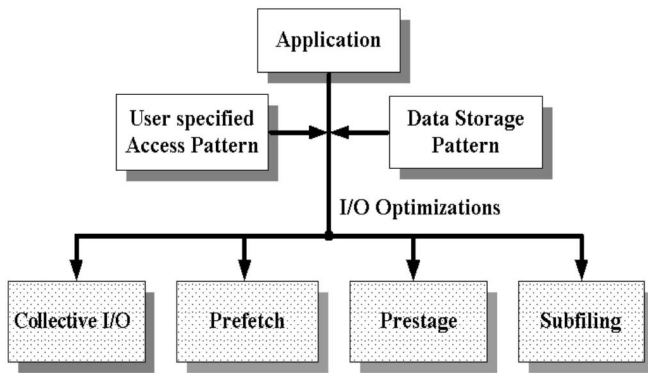


Fig. 4. Selecting an I/O optimization.

achieve better performance. Our library also provides other I/O optimization methods that are not found in MPI-I/O, but can be built on top of MPI-I/O using the access pattern information such as data prefetching (from disk or tape to memory), data prestaging (from tape to disk), and subfiling (for tape-resident data) [26]. For example, when the user is going to access a sequence of data sets and perform some computation on them sequentially, our library can overlap the I/O access and computation by prefetching or prestaging the next data set while the computation on the current data set continues. As another example, if the user will access a small chunk of data from a large tape-resident data set, our tape library, APRIL [26], will be called to achieve low latency in tape accesses. Another feature of the MDMS is that we provide mechanisms to locate the data by data set names, such as temperature or pressure rather than using file name and offset. The user can also query the MDMS to locate data sets in which she has particular interest and to devise application-specific access strategies for these data sets. Fig. 4 depicts a sketch of how an I/O optimization decision is made. In short, comparing the access pattern and storage pattern, and having access to the information about the location of the data set in the storage hierarchy, the MDMS can decide a suitable I/O optimization.

Note that, in our environment, the users' task is to convey the access pattern information to the MDMS and let the MDMS select a suitable I/O strategy for her. In addition to interprocessor access pattern information (hint), the MDMS also accepts information about, for example, whether the data set will be accessed sequentially, whether it is read-only for the entire duration of the program, and whether it will be accessed only once or repeatedly. An important problem now is in what part of the program the user should convey this information (hints). While one might think that such user-specified *hints* should be placed at the earliest point in the program to give the MDMS maximum time to develop a corresponding I/O optimization strategy, this may also hurt performance. For example, in receiving a hint, the MDMS can choose to act upon it, an activity that may lead to suboptimal I/O strategies had we considered the next hint. Therefore, sometimes *delaying hints* and issuing them to MDMS collectively might be a better choice. Of course, only the correlated hints must be issued together. While passing (access pattern) hints to file systems and runtime systems was proposed by other researchers [27], [24], [29], we believe that this is the first study that considers a large spectrum (variety) of performance-oriented hints in a unified framework.

The functions used by the MDMS to manipulate the database tables are given in Table 2. Fig. 5, on the other hand, shows a typical flow of calls using the MDMS. These routines are meant to substitute the traditional Unix I/O functions or MPI-I/O calls that may be used by the programmers when they want to read or dump data. They look very similar to typical Unix I/O functions in appearance, so the users do not have to change their programming practices radically to take advantage of state-of-the-art I/O optimizations. The flow of these functions can be described as follows:

1. **Initialization.** The MDMS flow starts with a call to the initialization() routine.
2. **Write.** The write operations start with create-association() that creates an association for the data sets that can be grouped together for manipulation. The create-association() returns an association-id that can be used later for collectively manipulating all the associated data sets. The subsequent function for the write operations is the save-initial() routine. This can be thought of as "file open" command in Unix-like I/O. Then, the user can use the save() function to perform data write operations to the storage hierarchy. Note that, in traditional Unix-like I/O, each data set needs a "file open," while in the MDMS library, there is only one "open:" the save-initial() routine collectively opens all the associated data sets. The write operations are ended with save-final() that corresponds to a "file close" operation in Unix-like I/O.
3. **Read.** The read operations start with the get-association() routine that obtains an association handle generated by the create-association() routine during a previous write operation. The next function to continue the read operations is load-initial() which, again, corresponds to "file-open" in Unix I/O. Then, the user can use the load() routine to perform read operations. The read operations are completed by the load-final() function. Note that the read and write operations can, of course, interleave.
4. **Finalization.** The MDMS flow is ended with the finalization() routine.

As stated earlier, the MDMS library provides transparent access to the database tables, thus users do not need to deal with these tables explicitly. The actions taken by the MDMS for a typical run session are as follows:

1. A row is added to the run table by set-run-table() to record the user-specified information about this run. Users can search this table by date and time to find information pertaining to a particular run.
2. For the data sets having similar characteristics such as the same dimension sizes, access pattern, etc., an association is created by create-association(). Each association with one or several data sets is inserted into the data set table. The access pattern table and storage pattern table are also accessed by the create-association(): The access pattern and storage pattern of each data set are inserted into these two tables, respectively. We expect the user to at least specify the access pattern for each data set. Note that,

TABLE 2  
Functions Used in the MDMS

Name	Functionality	Important Parameters	Tables Involved
initialization()	Initializes the MDMS environment	Application name	Application table
create-association()	Creates an association for the datasets with same behavior	Dataset name, access pattern	Dataset table
get-association()	Obtains the association for the datasets	Dataset name, access pattern	Dataset table
set-run-table()	Adds a row in the run table		Run table
load-initial()	Determines the file name and offset of the dataset; Opens the file; Determines I/O optimization method	Dataset, Association handle	Execution table, Access pattern table, Storage pattern table
load()	Determines whether prefetching should be performed; Performs I/O (read)	Association handle	None
load-final()	Closes the files involved	Association handle	None
save-initial()	Generates file names; Opens files for write; Determines I/O optimization method such as collective I/O, data sieving	Association handle	Execution table, Access pattern table, Storage pattern table
save()	Writes dataset	Dataset, Association handle	None
save-final()	Closes the files involved	Association handle	None

depending on the program structure, a data set might have multiple access patterns in different parts of the code. The MDMS also accepts user-specified storage pattern hints. If no storage pattern hint is given, the MDMS selects row-major layout (for C programs) or column-major layout (for Fortran programs).

- In load-init(), the file names, offsets, iteration number, etc., of a particular data set are searched from the execution table.
- In save-init(), the execution table may be searched to find out the file name for check-pointing. In save(), a row is inserted into execution table to record the current I/O activity.
- Steps 3 and 4 are repeated until the main loop where the I/O activity occurs is finished.

### 3 HIERARCHICAL STORAGE SYSTEM

The data sets that are generated by large-scale scientific applications might be too large to be held on the secondary storage devices permanently: thus, they have to be stored on tertiary storage devices (e.g., robotic tape) depending on their access profile. In many tape-based storage systems, the access granularity is a whole file [39]. Consequently, even if the program tries to access only a section of the tape-resident file, the entire file must be transferred from the tape to the upper level storage media (e.g., magnetic disk). This can result in poor I/O performance for many access patterns. The main optimization schemes in the MDMS we

have presented so far, such as collective I/O, prefetching and prestaging, could not help much when the user accesses only a small portion in a huge tape-resident data set as the tape access times would dominate. In this section, we present an optimization technique called *subfiling* that

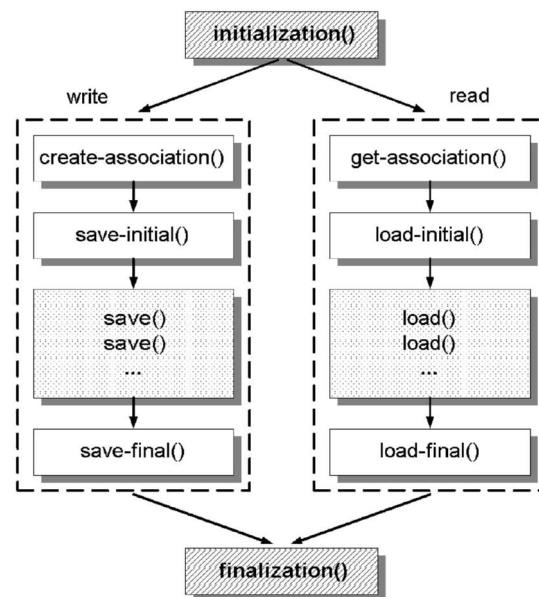


Fig. 5. A typical MDMS execution flow.




Fig. 6. (a) Interaction between the library calls, MPI-I/O, and HPSS. (b) Prefetching, prestaging, and migration.

can significantly reduce the I/O latencies in accessing tape-resident data sets.

### 3.1 Subfiling

We have developed and integrated into the MDMS a parallel runtime library (called APRIL) for accessing tape-resident data sets efficiently. At the heart of the library lies an optimization scheme called *subfiling*. In subfiling, instead of storing each tape-resident data set as a single large file, we store it as a collection of small subfiles. In other words, the original large data set is divided into uniform *chunks*, each of which can be stored independently in the storage hierarchy as a subfile. This storage strategy, however, is completely transparent to the user who might assume that the data set is stored in a single (logical) file. For read or write operations to the tape-resident data set, the start and end coordinates should be supplied by the user. The MDMS, in turn, determines the set of subfiles that collectively contain the required data segment delimited by the start and end coordinates. These subfiles are brought (using the APRIL API) from the tape to the appropriate storage device and the required data segment is extracted from them and returned to the user buffer supplied in the I/O call. The programmer is not aware of the subfiles used to satisfy the request. This provides a low-overhead (almost) random access for the tape-resident data with an easy-to-use interface.

The interaction between the library calls and the I/O software layers is depicted in Fig. 6a. Our current access to a storage hierarchy that involves tape devices is through HPSS (High Performance Storage System) [14]. The required subfiles are transferred (in a user-transparent manner) using the HPSS calls from the tape device to the disk device and then our usual MDMS calls (built on top of MPI-I/O) are used to extract the required subregions from each subfile. Fig. 6b shows some of the potential I/O optimizations between different layers.

### 3.2 Experiments with APRIL

We have conducted several experiments using the APRIL library API from within the MDMS. During our experiments, we have used the HPSS at the San Diego Supercomputing Center (SDSC). We have used the low-level

TABLE 3  
Access Patterns Used in the Experiments

Access Pattern	Start Coordinate	End Coordinate	Num of Floating Point Elements
A	(0, 0)	(1000, 1000)	$1 * 10^6$
B	(0, 0)	(4000, 1000)	$4 * 10^6$
C	(0, 0)	(24000, 1000)	$24 * 10^6$
D	(5000, 5000)	(6000, 6000)	$1 * 10^6$
E	(0, 0)	(80, 50000)	$4 * 10^6$
F	(0, 0)	(1000, 4000)	$4 * 10^6$
G	(6000, 6000)	(8000, 8000)	$4 * 10^6$

Each access pattern is delimited by a start coordinate and an end coordinate and contains all the data points in the rectangular region.

routines of the SDSC Storage Resource Broker (SRB) [2] to access the HPSS files. Table 3 shows the access patterns that we have experimented with (A through H). It also gives the start and end coordinates of the access patterns as well as the total number of elements requested by each access. In all these experiments, the global file was a two-dimensional matrix with  $50,000 \times 50,000$  floating-point elements. The chunk (subfile) size was set to  $2,000 \times 2,000$  (small chunks) and  $4,000 \times 4,000$  (large chunks) floating-point elements.

The results from our experiments are summarized in Table 4. The table gives the response times (in seconds) of the naive scheme (i.e., without subfiling) and the percentage gains achieved by our library using two subfile sizes (as given above) over the naive scheme. The results show that the library can, in general, bring about substantial improvements over the naive scheme for both read and write operations. The performance degradation in pattern C is due to the fact that subfiling causes a large number of small I/O calls while the naive approach completes the whole access with a single I/O call. We plan to eliminate these problems by developing techniques that help to select optimal subfile shapes given a set of potential access patterns. Smaller chunk sizes have the advantage of increased granularity. With the increased granularity the amount of redundant data transfer reduces, possibly increasing the efficiency. However, if the chunk size is too small, the number of files that has to be accessed increases, possibly reducing the overall efficiency. Hence, there is an optimal chunk size for each access pattern. If there is a single known pattern with fixed form accesses (i.e., accesses to the file has exactly the same form and size even if they are to different positions in the file), setting the chunk size to be equal to the access pattern will naturally result in the minimal delay. On the other hand, optimal chunk form for arbitrary access patterns or for a set of access patterns depends on the relative costs of transferring redundant data and accessing a file. Hence, it is related to the specific costs of the underlying system, such as data transfer rates and file access delay. Our initial observation is that the techniques proposed by Sarawagi [31] might be quite useful for this problem. Nevertheless, the results in Table 4 show that even a fixed chunk form increases the performance of the system significantly for almost all access patterns.

TABLE 4  
Execution Times and Percentage Gains for Write and Read Operations

Access Pattern	Write Operations			Read Operations		
	Times w/o chunking	Small Chunk Gain (%)	Large Chunk Gain (%)	Times w/o chunking	Small Chunk Gain (%)	Large Chunk Gain (%)
A	2774.0	96.1	94.5	784.7	85.2	77.1
B	2805.9	83.8	84.9	810.1	43.2	55.6
C	2960.3	8.8	37.9	793.3	-240.5	-172.4
D	3321.2	96.7	95.4	798.4	84.1	79.7
E	138723.3	96.0	97.2	39214.1	85.9	88.5
F	11096.3	95.9	96.4	3242.9	88.3	88.6
G	5095.2	91.2	96.5	1612.9	76.6	89.9

The second and fifth columns give the times for the naive I/O (without subfiling) in seconds. The remaining columns (except the first one) show the percentage improvements over the naive I/O method when subfiling is used.

#### 4 DESIGN OF THE INTEGRATED JAVA GRAPHICAL USER INTERFACE

As it is distributed in nature, our application development environment involves multiple resources across distant sites. For example, let us consider our current working environment that consists of different platforms and tools. We do program development using local HP or SUN workstations, the visualization tools used are installed on a Linux machine, our MDMS is located on another machine, and our parallel applications currently run on a 16-node IBM SP-2 distributed-memory message-passing architecture. Although these machines are within our department, they could be distributed across different locations in the Internet.

When the user starts to work on such a distributed environment without the help of our application development system, she normally needs to go through several steps that can be summarized as follows:

1. Log on to IBM SP2 and submit the parallel application.
2. When the execution of the application is complete, log on to the database host and use native SQL dialect to find the data set that would be needed for visualization.
3. Once the required data set has been found, transfer the associated file(s) manually, for example, using ftp, from SP2 (where data are located) to the visualization host (where visualization tools reside).
4. Log on to the visualization host (Linux machine) and start the visualization process.
5. Repeat Steps 2, 3, and 4 as long as there exist data sets 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 (IJ-GUI) is implemented and integrated to our application development

environment. The goal of the IJ-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 it is becoming a major language in distributed systems and it is easy to integrate Java in a Web-based environment. Java also provides the tools for a complete framework that addresses all aspects of managing the process of application development: processes and threads, database access, networking, and portability. In this environment, the users need to work only with IJ-GUI locally, rather than go to different sites to submit parallel applications or to do file transfers explicitly. Fig. 7 shows how IJ-GUI is related to other parts of our system. It actively interacts with three major parts of our system: with parallel machines to launch parallel applications, with the MDMS through JDBC to help users query metadata from databases, and with visualization tools. The main functions that IJ-GUI provides can be summarized as follows:

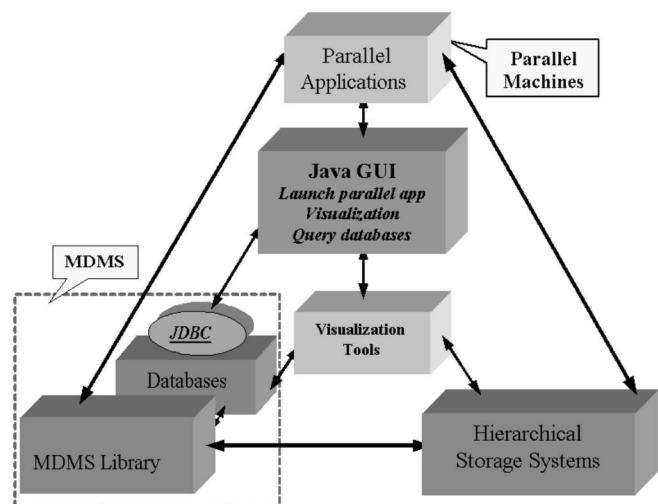


Fig. 7. Java GUI and the overall system.



- **Registering new applications.** To start a new application, the user needs to create a new suite of tables for the new application. Using the IJ-GUI, the user needs only to specify attributes (fields) of the run table and all other tables (e.g., storage pattern table, execution table, etc.) will be created *automatically* using the information provided in the run table.
- **Running applications remotely.** The applications typically run on some form of parallel architecture such as IBM SP2 that can be specified by the user when she registers a new application. Therefore, a remote shell command is used in IJ-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.
- **Data analysis and visualization.** Users can also carry out data analysis and visualization through the IJ-GUI. In general, data analysis is very application-specific and may come in a variety of flavors. For some applications, data analysis may simply calculate the maximum, minimum, or average value of a given data set whereas, for some others, it may be plugged into the application and calculate the difference between two data sets and decide whether the data sets should be dumped now or later. The current approach to the data analysis in our environment is to calculate the maximum, minimum, and arithmetic means of each data set generated. From the IJ-GUIs point of view, this process is no different than submitting a remote job. Visualization, on the other hand, is an important tool in large-scale scientific simulation, helping the users to inspect the inside nature of data sets. It is, in general, slightly more complicated than data analysis. First of all, the users' interests in a particular data set may be very arbitrary. Our approach is to list all the candidate data sets by searching the database using the user-specified characteristics such as maximum, minimum, means, iteration numbers, pattern, mode, etc. Then, the candidates are presented in a radio box for the user so that she can select the data set she wants. Second, the data sets are created by parallel machines, and they are located on parallel machines or stored in hierarchical storage systems. But, our visualization tools are installed in different locations. Therefore, inside IJ-GUI, we transparently copy the data from the remote parallel machine or hierarchical storage systems to the visualization host and then start the visualization process. The user does not need to check the MDMS tables explicitly for interesting data sets or perform data transfers manually. The only thing that she needs to do is to check-mark the radio box for interesting data sets, select a visualization tool (vtk, xv, etc.) and, finally, click the visualization button to start the process. The current visualization tools supported in our environment include Visualization Toolkit (vtk), Java 3D, and xv. Fig. 8 shows how the user visualizes data sets through vtk and xv.
- **Table browsing and searching.** Advanced users may want to search the MDMS tables to find the data sets of particular interest. Therefore, the table browsing and searching functions are provided in

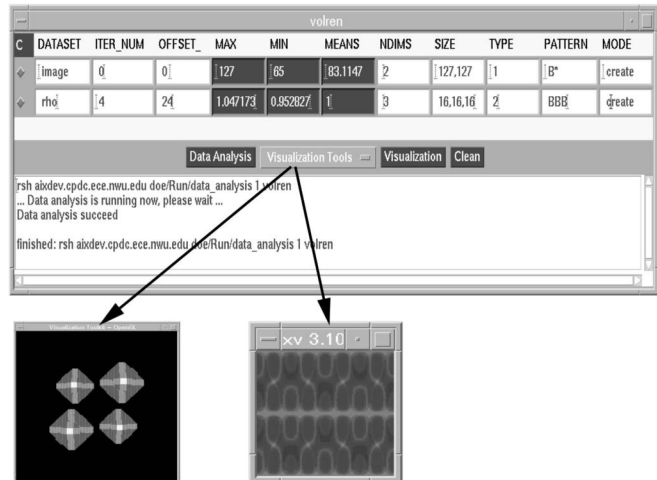


Fig. 8. A visualization example. The upper window shows the data sets along with their characteristics such as data sizes, iteration number (in which they are dumped), offset, pattern, etc. These data sets are chosen by the user for visualization. The lower windows show the visualization results for two different data sets, each using a different visualization tool.

the IJ-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.

- **Automatic code generator.** Our IJ-GUI relieves users great burden of working in a distributed system with multiple resources. For an application that has already been developed, the user would find it very easy to run her application with any parameters she wants: She can also easily carries out data analysis and visualization, search the database, and browse the tables. For a new application to be developed, however, although our high-level MDMS API is easy to learn and use, the user may need to make some efforts to deal with data structure, memory allocations, and argument selections for the MDMS functions. Although these tasks may be considered routine, we also want to reduce them to almost zero by designing an Automatic Code Generator (ACG) for MDMS API. The idea is that, given a specific MDMS function and other high-level information such as the access pattern of a data set, ACG will automatically generate a code segment that includes variable declarations, memory allocations, variable assignments, and identifications of as many of the arguments of that API as possible. The most significant feature of ACG is that it does not just work like a MACRO which is substituted for real codes: It may also consult databases for advanced information if necessary. For example, to generate a code segment for set-run-table(), which is to insert one row into the run table to record this run with user-specified attributes, our ACG would first search the database and return these attributes, then, it uses these attributes to fill out a predefined data structure as an argument in function set-run-table(). Without consulting the database, the user has to deal with these attributes by hand. Our ACG is integrated within our IJ-GUI as part of its functions. The user can simply copy the code segment generated and paste them in her own program.

TABLE 5  
Total I/O Times (in Seconds) for *astro2d*, *astro3d*, and Unstructured Applications  
(Data Set Size is 256 MB, 8 MB, and 64 MB, Respectively)

	astro2d		astro3d		unstructured	
	32 procs	64 procs	32 procs	64 procs	32 procs	64 procs
Original	23.46	39.67	109.93	211.47	547.61	488.13
Optimized	14.05	11.23	3.33	3.51	1.25	2.13

Currently, the IJ-GUI is implemented as a stand-alone system, we are in the process of embedding it into the Web environment, so the user can work in our integrated environment through a Web browser.

## 5 EXPERIMENTS

In previous sections, we have shown that our MDMS can provide an easy-to-use environment by relieving users from details of data storage architectures and file layouts. Unlike other systems that have to sacrifice significant performance in exchange for ease-of-use, our MDMS can still take advantage of state of the art I/O techniques to achieve performance improvements. In this section, we present some performance numbers from our current MDMS and IJ-GUI implementations. The experiments were run on an IBM SP-2 at Argonne National Lab. Each node of the SP-2 is RS/6000 Model 390 processor with 256 megabytes memory and has an I/O subsystem containing four nine gigabytes SSA disks attached to it.

We used four different applications: three of them are used to measure the benefits of collective I/O for disk-resident data sets; the last one is used to see how prestaging (i.e., staging data from tape to disk before they are needed) performs for tape-resident data and how prefetching (i.e., fetching data from disk to memory before they are needed) performs data already on disks. The current implementation of the APRIL library uses HPSS [14] as its main HSS interface to tape devices. HPSS is a scalable, next-generation storage system that provides standard interfaces (including an API) for communication between parallel processors and mass storage devices. Its architecture is based on the IEEE Mass Storage Reference Model Version 5 [13]. Through its parallel storage support by data striping, HPSS can scale upward as additional storage devices are added.

Table 5 shows the total I/O times for a two-dimensional astrophysics template (*astro2d*), a three-dimensional astrophysics code (*astro3d*), and for an unstructured (irregular data access pattern) code on the IBM SP-2. Here, Original

refers to the code without collective I/O and Optimized denotes the code with collective I/O. In all cases, the MDMS is run at Northwestern University. The important point here is that, in both the Original and the Optimized versions, the user code is essentially the same; the only difference is that the Optimized version contains access pattern hints and I/O read/write calls to the MDMS. The MDMS automatically determines that, for the best performance, collective I/O needs to be performed. As a result, impressive reductions in I/O times are observed. Since the number of I/O nodes are fixed on the SP-2, increasing the number of processors may cause (for some codes) an increase in the I/O time.

Note that an experienced programmer who is familiar with file layouts and storage architectures can obtain the same results by manually optimizing these three applications using collective I/O. This requires, however, significant programming time and effort on the programmers' part. Our work and results show that such improvements can also be possible using a smart metadata management system and requiring users to indicate only access pattern information.

Our next example (Table 6) is a parallel volume rendering application (*volren*). As in previous experiments, the MDMS is run at Northwestern University. The application itself, on the other hand, is executed at Argonne National Lab's SP-2 and the HPSS at the San Diego Supercomputer Center (SDSC) is used as the HSS. In the Original code, four data files are opened and parallel volume rendering is performed. In the Optimized code, the four data sets (corresponding to four data files) are associated with each other, and prestaging (from tape to disk) is applied for these data sets. Table 6 gives the total read times for each of the four files for the Original and Optimized codes for four and eight processor cases, respectively. The results reveal that, for both four and eight processor cases, prestaging reduces the I/O times significantly. We need to mention that, in every application we experimented with in our environment, the time spent by

TABLE 6  
Total I/O Times (in Seconds) for *volren* on Four and Eight Processors (Data Set Size is 64 MB) and Average I/O Times

File No →	volren 4 procs				volren 8 procs			
	1	2	3	4	1	2	3	4
Original	31.18	19.20	61.86	40.22	18.79	37.69	21.02	14.70
Optimized	11.90	11.74	20.10	18.38	10.74	6.23	4.49	6.42

TABLE 7  
Average I/O Times (in Seconds) for *volren*  
(Data Set Size is 2 MB)

	4 procs	8 procs
Original	2.27	1.34
Optimized	1.91	1.15

the application in negotiating with the MDMS was less than one second. When considering the significant runtime improvements provided by I/O optimizations, we believe that this overhead is not great.

Finally, we also measure the benefits of prefetching in *volren*. We assume the data sets are stored on local SP-2 disks. In the `Original` code, four data files are opened and computations are performed sequentially. In the `Optimized` code, prefetching (from disk to memory) is applied to the next data file when each processor is doing computation on the current data file. Consequently, the I/O time and computation time are overlapped. Table 7 shows the average read times for the four files for the `Original` and `Optimized` codes for four and eight processor cases, respectively. The results demonstrate that, for both four and eight processor cases, prefetching decreases the I/O time by 15 percent. Actually, prefetching and prestaging are complementary optimizations. Our environment is able to take advantage of overlapping prestaging, prefetching, and computation, thereby maximizing the I/O performance.

## 6 RELATED WORK

Numerous techniques for optimizing I/O accesses have been proposed in the literature. These techniques can be classified into three categories: the parallel file system and runtime system optimizations [22], [8], [10], [19], [21], [16], compiler optimizations [4], [20], [17], and application analysis and optimization [20], [6], [29], [17], [7], [38]. Brown et al. [5] proposed a metadata system on top of HPSS using DB2 DBMS. Our work, in contrast, focuses more on utilizing state-of-the-art I/O optimizations with minimal programming effort. Additionally, the design flexibility of our system allows us to easily experiment with other hierarchical storage systems as well. The use of high-level unified interfaces to data stored on file systems and DBMS is investigated by Baru et al. [2]. Their system maintains metadata for data sets, resources, users, and methods (access functions), and provides the ability to create, update, store, and query this metadata. While the type of metadata maintained by them is an extension of metadata maintained by a typical operating system, our metadata involves performance-related metadata as well which enables automatic high-level I/O optimizations as explained in this paper.

## 7 CONCLUSIONS

This paper has presented a novel application development environment for large-scale scientific computations. At the

core of our framework is the Metadata Database Management System (MDMS) framework, which uses relational database technology in a novel way to support the computational science analysis cycle described at the beginning of this paper in Fig. 1. A unique feature of our MDMS is that it relieves users from choosing best I/O optimizations such as collective I/O, prefetching, prestaging, etc., that may typically exceed the capabilities of a computational scientist who manipulates large data sets. The MDMS itself is made useful by the presence of a C application programming interface (API) as well as an integrated Java Graphical User Interface (IJ-GUI), which eliminates the need for computational scientists to work with complex database programming interfaces such as SQL and its embedded forms, which typically varies from vendor to vendor. The IJ-GUI itself is a key component of the system that allows us to transparently make use of heterogeneously distributed resources without regard to platform. We also presented an optimization for tape-resident data sets, called subfiling, that aims at minimizing the I/O latencies during data transfers between secondary storage and tertiary storage. Our performance results demonstrated that our novel programming environment provided both ease-of-use and high-performance.

We are currently investigating other tape-related optimizations and trying to fully-integrate MDMS with hierarchical storage systems such as HPSS. We are also examining other optimizations that can be utilized in our distributed environment when the user carries out visualization. Regarding future tertiary storage media, CD-ROMS are a good competitor with tape-based storage systems for large-scale applications. Overall, the work presented in this paper is a first attempt to unify the best characteristics of databases, parallel file systems, hierarchical storage systems, Java, and the Web, to enable effective high-level data management in scientific computations.

## ACKNOWLEDGMENTS

This research was supported by the Department of Energy under the Accelerated Strategic Computing Initiative (ASCI) Academic Strategic Alliance Program (ASAP) Level 2, under subcontract no. W-7405-ENG-48 from Lawrence Livermore National Laboratories. The authors would like to thank Reagan Moore for discussions and help with SDSC resources. They would also like to thank Mike Wan and Mike Gleicher of SDSC for helping them with the implementation of the volume rendering code and in understanding the SRB and the HPSS. They thank Larry Schoof and Wilbur Johnson for providing the unstructured code used in this paper. They also thank Rick Stevens and Rajeev Thakur of ANL for various discussions on the problem of data management. They also thank Jaechun No for her help with the astrophysics application used in the experiments. Finally, they would like to thank Celeste Matarazzo, John Ambrosiano, and Steve Louis for discussions and their input.

## REFERENCES

- [1] C. Baru, R. Frost, J. Lopez, R. Marciano, R. Moore, A. Rajasekar, and M. Wan, "Meta-Data Design for a Massive Data Analysis System," *Proc. CASCON Conf.*, 1996.
- [2] C. Baru, R. Moore, A. Rajasekar, and M. Wan, "The SDSC Storage Resource Broker," *Proc. CASCON Conf.*, Dec. 1998.
- [3] R. Bennett, K. Bryant, A. Sussman, R. Das, and J.S. Jovian, "A Framework for Optimizing Parallel I/O," *Proc. Scalable Parallel Libraries Conf.*, 1994.
- [4] R. Bordawekar, A. Choudhary, K. Kennedy, C. Koelbel, and M. Paleczny, "A Model and Compilation Strategy for Out-of-Core Data Parallel Programs," *Proc. ACM Symp. Principles and Practice of Parallel Programming*, pp. 1-10, 1995.
- [5] P. Brown, R. Troy, D. Fisher, S. Louis, J.R. McGraw, and R. Musick, "Meta-Data Sharing for Balanced Performance," *Proc. First IEEE Meta-Data Conf.*, 1996.
- [6] P. Cao, E. Felten, and K. Li, "Application-Controlled File Caching Policies," *Proc. Summer USENIX Technical Conf.*, pp. 171-182, 1994.
- [7] C.-M. Chen and R.K. Sinha, "Analysis and Comparison of Decustering Schemes for Interactive Navigation Queries," *IEEE Trans. Knowledge and Data Eng.*, vol. 12, no. 5, pp. 763-778, Sept./Oct. 2000.
- [8] 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, 1994.
- [9] P. Corbett, D. Feitelson, J.-P. Prost, G. Almasi, S.J. Baylor, A. Bolmarich, Y. Hsu, J. Satran, M. Snir, R. Colao, B. Herr, J. Kavaky, T. Morgan, and A. Zlotek, "Parallel File Systems for the IBM SP Computers," *IBM Systems J.*, vol. 34, no. 2, pp. 222-248, Jan. 1995.
- [10] P. Corbett, D. Feitelson, S. Fineberg, Y. Hsu, B. Nitzberg, J. Prost, M. Snir, B. Traversat, and P. Wong, "Overview of the MPI-IO Parallel I/O Interface," *Proc. Third Workshop I/O in Parallel and Distributed Systems*, 1995.
- [11] P.F. Corbett, D.G. Feitelson, J.-P. Prost, and S.J. Baylor, "Parallel Access to Files in the Vesta File System," *Proc. Supercomputing Conf.*, pp. 472-481, 1993.
- [12] T.H. Cormen and D.M. Nicol, "Out-of-Core FFTS with Parallel Disks," *ACM SIGMETRICS Performance Evaluation Rev.*, pp. 3-12, 1997.
- [13] R.A. Coyne and H. Hulen, "An Introduction to the Mass Storage System Reference Model," *Proc. 12th IEEE Symp. Mass Storage Systems*, 1993.
- [14] R.A. Coyne, H. Hulen, and R. Watson, "The High Performance Storage System," *Proc. Supercomputing Conf.*, 1993.
- [15] J.R. Davis, "Datalinks: Managing External Data with db2 Universal Database," *IBM Corporation White Paper*, 1997.
- [16] J. del Rosario, R. Bordawekar, and A. Choudhary, "Improved Parallel I/O via a Two-Phase Run-Time Access Strategy," *Proc. IPPS Workshop Input/Output in Parallel Computer Systems*, 1993.
- [17] J. del Rosario and A. Choudhary, "High Performance I/O for Parallel Computers: Problems and Prospects," *IEEE Computer*, Mar. 1994.
- [18] M. Drewry, H. Conover, S. McCoy, and S. Graves, "Meta-Data: Quality vs. Quantity," *Proc. Second IEEE Meta-Data Conf.*, 1997.
- [19] C.S. Ellis and D. Kotz, "Prefetching in File Systems for MIMD Multiprocessors," *Proc. Int'l Conf. Parallel Processing*, pp. 306-314, 1989.
- [20] M. Kandaswamy, M. Kandemir, A. Choudhary, and D. Bernholdt, "Performance Implications of Architectural and Software Techniques on I/O-Intensive Applications," *Proc. Int'l Conf. Parallel Processing*, 1998.
- [21] J.F. Karpovich, A.S. Grimshaw, and J.C. French, "Extensible File Systems (ELFS): An Object-Oriented Approach to High Performance File I/O," *Proc. Ninth Ann. Conf. Object-Oriented Programming Systems, Languages, and Applications*, pp. 191-204, 1994.
- [22] D. Kotz, "Multiprocessor File System Interfaces," *Proc. Second Int'l Conf. Parallel and Distributed Information Systems*, pp. 194-201, 1993.
- [23] D. Kotz, "Disk-Directed I/O for MIMD Multiprocessors," *Proc. Symp. Operating Systems Design and Implementation*, pp. 61-74, 1994.
- [24] T. Madhyastha and D. Reed, "Intelligent, Adaptive File System Policy Selection," *Proc. Conf. Frontiers of Massively Parallel Computing*, pp. 172-179, 1996.
- [25] MCAT, <http://www.npaci.edu/DICE/SRB/mcat.html>, 2003.
- [26] G. Memik, M. Kandemir, A. Choudhary, and V.E. Taylor, "APRIL: A Run-Time Library for Tape Resident Data," *Proc. 17th IEEE Symp. Mass Storage Systems*, 2000.
- [27] T.C. Mowry, A.K. Demke, and O. Krieger, "Automatic Compiler-Inserted I/O Prefetching for Out-of-Core Applications," *Proc. Second Symp. Operating Systems Design and Implementation*, pp. 3-17, 1996.
- [28] J. Newton, "Application of Meta-Data Standards," *Proc. First IEEE Meta-Data Conf.*, 1996.
- [29] R.H. Patterson, G.A. Gibson, and M. Satyanarayanan, "A Status Report on Research in Transparent Informed Prefetching," *ACM Operating Systems Rev.*, pp. 21-34, 1993.
- [30] B. Rullman, "Paragon Parallel File System," External Product Specification, Intel Supercomputer Systems Division, 1996.
- [31] S. Sarawagi, "Query Processing in Tertiary Memory Databases," *Proc. 21st Very Large Databases Conf.*, 1995.
- [32] M. Stonebraker, *Object-Relational DBMSs: Tracking the Next Great Wave*. Morgan Kaufman, 1998.
- [33] M. Stonebraker and L.A. Rowe, "The Design of Postgres," *Proc. ACM SIGMOD Int'l Conf. Management of Data*, pp. 340-355, 1986.
- [34] R. Thakur, W. Gropp, and E. Lusk, "A Case for Using MPI's Derived Datatypes to Improve I/O Performance," *Proc. Supercomputing Conf.: High Performance Networking and Computing*, 1998.
- [35] R. Thakur, W. Gropp, and E. Lusk, "On Implementing MPI-IO Portably and with High Performance," Preprint ANL/MCS-P732-1098, Math. and Computer Science Division, Argonne Nat'l Laboratory, 1998.
- [36] R. Thakur, W. Gropp, and E. Lusk, "Data Sieving and Collective I/O in Romio," *Proc. Seventh Symp. Frontiers of Massively Parallel Computation*, 1999.
- [37] S. Toledo and F.G. Gustavson, "The Design and Implementation of Solar, a Portable Library for Scalable Out-of-Core Linear Algebra Computations," *Proc. Fourth Ann. Workshop I/O in Parallel and Distributed Systems*, 1996.
- [38] P. Triantafillou and T. Papadakis, "Continuous Data Block Placement and Elevation from Tertiary Storage in Hierarchical Storage Servers," *Cluster Computing: The J. Networks, Software Tools, and Applications*, 2001.
- [39] UniTree User Guide, Release 2.0, UniTree Software, Inc., 1998.



Xiaohui Shen received the BS and MS degrees from Department of Computer Science at Tsinghua University, Beijing, and the PhD degree from Northwestern University, Illinois, in computer engineering in January 2001. He is currently a senior software engineer at Core Technology Department, Personal Communication Sector, Motorola Inc. His research interests include parallel and distributed computing, data management, I/O and storage systems, and mobile computing. He has served as a program committee member for several international conferences and workshops on high-performance and mobile computing.



Wei-keng Liao received the PhD degree in computer and information science from Syracuse University in 1999. He is a research assistant professor in the Electrical and Computer Engineering Department at Northwestern University. His research interests are in the area of high-performance computing, parallel I/O, data mining, and data management for large-scale scientific applications.



**Alok Choudhary** received the PhD degree from the University of Illinois, Urbana-Champaign, in electrical and computer engineering in 1989, and the MS degree from the University of Massachusetts, Amherst, in 1986. Since September 2000, he has been a professor in the Electrical and Computer Engineering Department and The Kellogg School of Management at Northwestern University. From 1989 to 1996, he was a faculty member in the Electrical and

Computer Engineering Department at Syracuse University. He received the US National Science Foundation's Young Investigator Award in 1993, an IEEE Engineering Foundation award, and an IBM Faculty Development award. His main research interests are in high-performance computing and communication systems and their applications in many domains including information processing (e.g., data mining) and scientific computing (e.g., scientific discoveries). In particular, his interests lie in the design and evaluation of architectures and software systems (from system software such as runtime systems to compilers), high-performance servers, high-performance databases, and input-output. He has published more than 250 papers in various journals and conferences in the above areas. He has also written a book and several book chapters on the above topics. He is a senior member of the IEEE.



**Gokhan Memik** received the BS degree in 1998 from Bogazici University, Istanbul, Turkey, and the MS degree in 2000 from Northwestern University, Evanston, Illinois, both in computer engineering, and the PhD degree in electrical engineering from the University of California, Los Angeles in 2003. He has been an assistant professor with the Electrical and Computer Engineering Department at Northwestern University since September 2003. His research interests are broadly in the area of computer architecture with an emphasis on high-performance embedded systems. He has done active research on network processors, general-purpose microarchitectures, storage systems, and design automation for embedded systems.



**Mahmut Kandemir** received the BSc and MSc degrees in control and computer engineering from Istanbul Technical University, Istanbul, Turkey, in 1988 and 1992, respectively. He received the PhD degree from Syracuse University, Syracuse, New York in electrical engineering and computer science in 1999. He has been an assistant professor in the Computer Science and Engineering Department at the Pennsylvania State University since August

1999. His main research interests are optimizing compilers, I/O intensive applications, and power-aware computing. He is a member of the IEEE and the ACM.

▷ **For more information on this or any other computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.**