

A Virtual File System Interface for Computational Grids

Abdulrahman Azab and Hein Meling

Dept. of Electrical Engineering and Computer Science
University of Stavanger, N-4036 Stavanger, Norway

Abstract. Developing applications for distributed computation is becoming increasingly popular with the advent of grid computing. However, developing applications for the various grid middleware environments require attaining intimate knowledge of specific development approaches, languages and frameworks. This makes it challenging for scientists and domain specialists to take advantage of grid frameworks. In this paper, we propose a different approach for scientists to gain programmatic access to the grid of their choice. The principle idea is to provide an abstraction layer by means of a *virtual file system* through which the grid can be accessed using well-known and standardized system level operations available from virtually all programming languages and operating systems. By abstracting away low-level grid details, domain scientists can more easily gain access to high-performance computing resources without learning the specifics of the grid middleware being used. We have implemented such a virtual file system on HIMAN, a peer-to-peer grid middleware platform. Our initial experimental evaluation shows that the virtual file system only cause a negligible overhead during task execution.

1 Introduction

Grid computing is an appealing concept to support the execution of computationally intensive tasks; grid computing generally refers to the coordination of the collective processing power of scattered computing nodes to accommodate such large computations. Peer-to-peer grid computing extends this idea, enabling all participants to provide and/or consume computational power in a decentralized fashion. Although appealing at first sight, grid computing frameworks are complex pieces of software and thus typically come with some obstacles limiting adoption to expert programmers capable of and willing to learn to program using the grid middleware application programming interfaces (APIs). Therefore, a major challenge faced by developers of grid computing frameworks is to expand the base of grid application developers to include non-expert programmers, e.g. domain specialists using their own domain-specific tools. One approach is to devise transparency mechanisms to hide the complexities of the grid computing system. In this paper, we propose the design of a programming interface for a peer-to-peer grid middleware through a *virtual file system* (VFS) [5, 11, 22]. The rationale behind this idea is that the file system interface is well-known, and is

accessible from just about any domain specific tool or even simple file-based commands. The approach also lends itself to devising a file system organization exposing varying degrees of complexity to its users. For example, complex management operations can be performed using one part of the file system hierarchy, whereas domain specialists can submit tasks through another part of the file system limiting the exposure to a most essential part of the system. The main difference between this approach and the existing approaches is that it uses the most familiar interface between the user and the OS, i.e. FS, which is a data management interface to provide interaction with a computation management environment, i.e. the Grid.

We have implemented such a VFS layer above a middleware core based on our HIMAN platform [1, 3]. The VFS layer sits between the end user/programmer and the client program enabling them to assign a task to be executed on the grid simply by calling the `write()` function of the virtual file system on a specific virtual path. Similarly, when task execution has completed, the user can easily obtain the results by invoking the `read()` function on another virtual path. To evaluate the overhead imposed by the VFS layer, we compared a simple parallel matrix multiplication task running with and without the VFS layer. The results confirm our intuition that the VFS layer only adds a negligible computational overhead.

The rest of the paper is organized as follows: Section 2 gives an overview of the related efforts. Section 3 presents a general overview of grid middleware architectures. Section 4 describes the new added virtual file system layer and how it interacts with the other components in the middleware. Section 5 presents and discusses the results obtained from the performed experiments. Section 6 presents conclusions and future work plans.

2 Related Work

The concept of virtual file systems was first introduced in the Plan 9 operating systems [22] to support a common access style to all types of IO devices, be it disks or network interfaces. This idea has since been adopted and included in the Linux kernel [11], and support is also available on many other platforms, including Mac OS X [18] and Windows [5, 8]. There are a wide range of applications of such a virtual file system interface, including sshfs, procs, and YouTubeFS.

In the context of grid computing, providing access to a (distributed) file system architecture is both useful and commonplace. Transmitting the data necessary for computation is made easier with a distributed file system, and a lot of research effort has gone into data management for the grid, e.g. BAD-FS [4], and finding efficient protocols for data transfer, e.g. GridFTP [14]. BAD-FS [4] is different from many other distributed file systems in that it gives explicit control to the application for certain policies concerning caching, consistency and replication. A scheduler is provided with BAD-FS that takes these policies into account for different workloads. The Globus XIO [2] provides a file system-based API for heterogeneous data transfer protocols aimed for the Grid. This allows

applications to take advantage of newer (potentially) more efficient data transfer protocols in the future, assuming they too provide the XIO APIs. These file systems tackle the problem of interacting with the data management component and file transfer protocols of a data grid environment. However, none of them try to solve the generic interaction with the computational grid middleware for task submission, results collection in a suitable manner.

Some work has been carried out in the trend of using VFS to provide transparent interaction with remote target systems in a distributed environment. UEM [27] is a unified execution model for cloud and cluster computing which provides VFS interface for the execution, monitoring, and control of remote processes. This is carried out by presenting each process in the system as a virtual directory. Within each process' directory, virtual files are serving as interfaces for information retrieval and control of the remote process. The structure of nodes within the distributed system is displayed as a hierarchical directory structure in away that each node directory includes a number of sub-directories one for each remotely called process within the node. The drawback of this approach is that the user has to specify at which node should the process be allocated. This can be simple in case of simple distributed structure of the system. But as the complexity of the system structure increases, the complexity of the VFS structure will also increase, which will in turn reduce the transparency of the interface. Parrot [25] is an interposition agent [15] which provides seamless integration between standard Unix applications and remote storage systems. This integration makes a remote storage system appear as a file system to a Unix application. One drawback is that the user has to specify the location of the remote machine. Besides, all commands has to be included in a Parrot command, which decreases the provided transparency. Many other systems [10, 17, 19, 20, 24] also provide an intermediate layer between local user interfaces and remote applications for providing transparent access to the target applications. All of these systems require that the user has to know how to deal with the front end application which is system specific.

Our proposed VFS layer provides an interface which is well known to all computer users and locates itself as the front end application. The user does not have to provide any information about the remote machine or to be aware of the structure of the distributed system. It also hides the complexity of the process management, i.e. scheduling and fault-tolerance, from the user. To our knowledge, most of the existing computational grid middleware environments [6, 7, 9, 13, 16, 17, 21, 23] require a special language level API for interaction with the scheduler and to collect the results.

3 HIMAN Overview

The VFS layer has been implemented on HIMAN [3], a pure peer-to-peer grid middleware that support serial and parallel computational tasks. HIMAN has three main components: (i) a *worker* component responsible for task execution, (ii) a *client* component responsible for task submission and execution monitoring,

and (iii) a *broker* responsible for the task allocation. All nodes in the grid have all three components. The client and broker of a submitting node is then responsible for task submission, allocation, and execution monitoring in a decentralized scheduling fashion. Any node can submit tasks, and can also serve as executor for other tasks at the same time. Since HIMAN enables the execution of both serial and parallel tasks [1], the current GUI interface requires the user to specify many configuration parameters such as: locations of the code and data files, number of subtasks, execution environment, and scheduling policy. This can be too complex for the non-expert users.

4 Virtualizing Grid Access

Existing grid computing platforms [6, 12, 13, 26] require that the user configure a wide range of parameters, have sufficient background in distributed computing, and proficient knowledge of the grid middleware APIs to be able to execute tasks on the grid. This might be appropriate for expert users, but for regular users, a simple and familiar interface is desirable. By far the most ubiquitous programming interface available on computer systems is the file system interface. Thus, our proposed VFS layer placed above our HIMAN grid computing middleware exposes such an interface to the application programmer, making it easy to interact with the middleware using a well-known interface. Moreover, it enables users to submit tasks using the `write()` function, and monitor and collect results from the computation using the `read()` function of the VFS layer.

4.1 The VFS Layer

Our VFS layer is implemented as an additional interface layer above the client using the Callback file system library [5]. Callback is similar to the FUSE library available on Linux and Mac OS X [11, 18], and enable building virtual file systems in user space. The Callback library contains a user mode API to communicate with user applications, and a kernel mode file system driver for communicating with the Windows kernel. The role of the Callback library in our VFS layer is to enable users to provide input files and collect result files using simple file system commands. In order to make the proposed VFS applicable and easy to implement in other grid systems which are based on different platforms, the communication between the VFS layer and the client is performed by means of simple UDP text-based commands. In order to implement VFS on another grid middleware, a small UDP communication module has to be added to the client.

The installation of the VFS layer on a grid client machine is very simple. The user simply runs an MSI package, which installs the VFS layer as a Windows service. The service is configured to startup automatically when the user logs on, or manually by executing a `net start` command. Once the service is started, the VFS drive will be mounted. The VFS layer will communicate with the grid client upon the execution of specific file system commands. The virtual drive will be unmounted automatically when the user logs off, or it can be unmounted manually by executing a `net stop` command.

The VFS layer is composed of two main modules: Task Submission Module (TSM) and Result Collection Module (RCM).

Task Submission Module. TSM contains two procedures. The first is the virtual volume procedure, and will be called when the Callback file system is mounted, and creates a virtual volume visible in Windows Explorer. This new volume will be used by the user to provide the input files. The second is the task submission procedure and is responsible for forwarding the input files to the client in order to start the task submission process. Task submission requests are handled in the `CbFsFsWrite()` function of the Callback user mode library, so that it is triggered only when a `write()` command for a specific file is executed on the virtual drive. The default case is when the program file is copied to the virtual drive. This can be changed by the user to determine which `write()` command will trigger the TSM. The task submission procedure depicted in Fig. 1 involves the following steps:

1. The user executes a `write()` command, through the file system interface (e.g. Windows Explorer), to write the input files into the virtual volume.
2. The file system interface forwards the command to the Windows kernel.
3. The kernel forwards the command to the Callback file system driver.
4. The file system driver responds by calling the `CbFsFsWrite()` function in the Callback user mode library.
5. This function invokes the task submission procedure in the TSM.
6. The task submission procedure responds by presenting the input files in the file list of the virtual volume, so that they can be accessed with the file system interface, and forwarding the input files to the client.
7. The client submits the task to the grid.
8. During the execution, the user can monitor the execution process (i.e. the progress) by executing a `read()` command on a specific text file on the virtual drive.

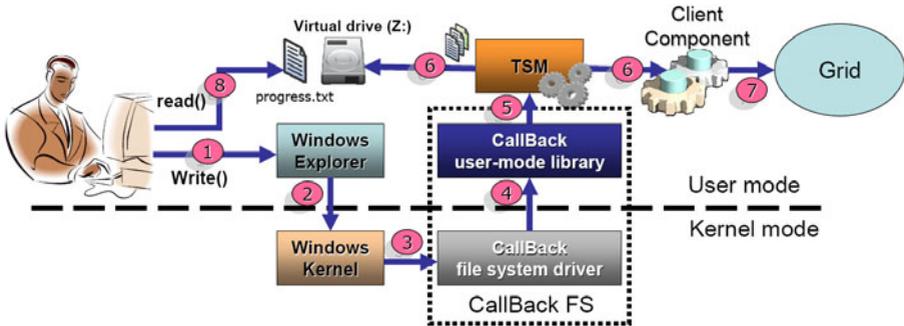


Fig. 1. Task submission procedure

Result Collection Module. The objectives of the result collection module are twofold: (i) provide a file system interface through which the user can collect the result files from task execution, and (ii) signal the completion of task execution to the user. The result collection procedure depicted in Fig. 2 involves the following steps:

1. When task execution is completed, the results are sent to the client process on the submitting node [3].
2. The client sends the collected results to the RCM. The result files are then stored in a temporary physical directory.
3. The RCM responds by creating a new virtual directory (e.g. **Results**) in the virtual drive mapping the temporary physical directory. Creation of the **Results** directory indicates to the user that task execution has complete.
4. The user can collect the result files by executing a **move** command on the files in the **Results** directory to move the result files to a physical path.
5. Windows Explorer forwards the command to the kernel.
6. The kernel forwards the command to the Callback file system driver.
7. The file system driver responds by calling the `CbFsRenameOrMoveEvent()` event procedure in the Callback user mode library.
8. The `CbFsRenameOrMoveEvent()` procedure responds by moving the result files to the given physical path, and removing them from the file list of the **Results** directory of the virtual drive.

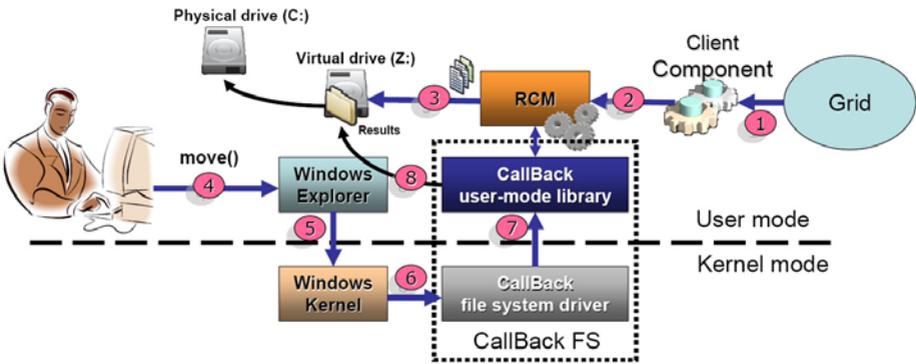


Fig. 2. Result collection procedure

4.2 A Simple Example

In this section we provide a simple example describing how to use the VFS layer for executing a parallel task of matrix multiplication. To accomplish this, the user must provide as input: an executable code file (e.g. `code.dll`), and the necessary data input files (e.g. `input?.data`). The number of input files correspond to the number of parallel subtasks that will be generated by the HIMAN middleware during execution of the task.

For illustration, the DOS command prompt is used (interactively) below to issue file system commands to interact with the grid middleware. The virtual drive is mounted on drive Z:, and the input files are stored in C:\Input\. The following steps describe the whole execution procedure, including results collection.

1. The user writes the data files to the virtual drive by typing:

```
copy C:\Input\*.data Z:
```

2. The user writes the code file to the virtual drive by typing:

```
copy C:\Input\code.dll Z:
```

As explained in Section 4.1, this command will trigger the TSM to invoke the client to begin task execution, and to create virtual file `progress.txt` on drive Z: for monitoring progress.

3. During the execution, the user can monitor the execution progress for the subtasks by typing:

```
type progress.txt
```

```
Subtask 1. Worker:193.227.50.201 Progress: 10%
Subtask 2. Worker:74.225.70.20 Progress: 15%
Subtask 3. Worker:87.27.40.100 Progress: 20%
Subtask 4. Worker:80.50.96.119 Progress: 6%
Subtask 5. Worker:211.54.88.200 Progress: 12%
```

4. The user can repeat the above command to keep up with the execution progress. The user need not consider management issues such as scheduling, fault tolerance, and connectivity. These issues are seamlessly handled by the HIMAN grid middleware [3]. In case of a worker failure, the worker address will be changed for the associated subtask in the progress file.
5. When the execution of one or more subtasks is completed, this will be revealed in the `progress.txt` file as follows:

```
Subtask 1. COMPLETED
Subtask 2. Worker:74.225.70.20 Progress: 90%
Subtask 3. Worker:87.27.40.100 Progress: 88%
Subtask 4. COMPLETED
Subtask 5. COMPLETED
```

6. Upon the completion of all subtasks, the `Results` virtual directory will appear, enabling the user to move all the virtual files to a physical directory as follows:

```
move Z:\Results\*.* C:\Results
```

5 Initial Performance Evaluation

In order to demonstrate the applicability of our VFS layer, we have performed a simple experimental evaluation to reveal the overhead caused by the VFS layer compared to interacting directly with the HIMAN middleware (through a GUI interface).

Two experiments were performed using a classic parallel matrix multiplication task for multiplying two square matrices of size: a) 1500×1500 and, b) 2100×2100 . In both cases, different number of workers (i.e. parallel subtasks) varying from one to six were used. The implemented benchmark is to calculate the task submission time (TST) for both cases: direct submission (using the GUI), and submission through the VFS interface. In case of direct submission, TST interval starts when the user submits the task after setting all required configuration options. In case of VFS, the TST starts from the first step in the TSM, Section 4.1. TST ends when all subtasks have been transferred to workers. Of course in case of direct submission, the user will take some time setting the configuration options, but as long as this time depends on the user it is not included in the TST. The result collection phase is not considered since in both cases the result files are copied to a physical path and the user has to collect them. The results are shown in Fig. 3.

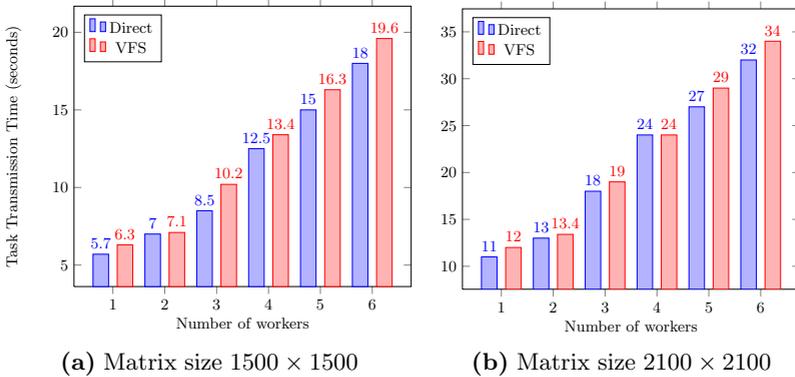


Fig. 3. Task transmission time against the number of workers using Direct and VFS interface submissions

In Fig. 3, it is clear that the computational overhead which is represented by the increase in the task transmission time in case of using the VFS interface instead of the built in client GUI, is nearly negligible. Since the transmission of the input files to workers is done in a serial fashion, the overhead is due to the time taken by the VFS to build the virtual files in memory and to communicate with the client for transmitting each of input files. The reason behind implementing Serial transmission in HIMAN is to avoid the large traffic in the submitting node which will be resulted from sending many data files at the same time, especially in case of large files.

The two approaches, direct and VFS, have identical CPU overhead simply because in both cases all pre-processing procedures are carried out by the grid client. The memory overhead is also identical since in both cases, the task input files are loaded into memory in order to be processed by the grid client upon the execution.

6 Conclusions and Future Work

Given the complexity of grid computing middleware, providing an easy to use interface for task submission is a significant challenge. In this paper, we have proposed a new technique for accessing computational grid middleware through a virtual file system interface. The proposed technique has been implemented on the HIMAN middleware, enabling non-expert users to submit compute tasks to the grid. Our initial evaluation indicate that the overhead imposed by the VFS interface over the direct interaction approach is negligible.

Our current implementation support serial and simple parallel task execution; in future work we will extend the VFS layer with support for complex parallel task execution. Moreover, we also plan to design a VFS-based grid portal that supports multiple grid computing frameworks, e.g. Globus and Condor. To accomplish this we will leverage the FUSE [11] framework on Linux to provide cross platform support for programming languages and runtime environments that can run on multiple operating systems. We will also add support for multi-user scenarios taking security and scalability issues into account.

References

1. El-Desoky, A.E., Ali, H.A., Azab, A.A.: A pure peer-to-peer desktop grid framework with efficient fault tolerance. In: ICCEs 2007, Cairo, Egypt, pp. 346–352 (2007)
2. Allcock, W., Bresnahan, J., Kettimuthu, R., Link, J.: The globus extensible input/output system (xio): A protocol independent io system for the grid. In: IPDPS 2005: Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2005) - Workshop 4, Washington, DC, USA, pp. 179.1. IEEE Computer Society Press, Los Alamitos (2005)
3. Azab, A.: HIMAN: A Pure Peer-to-Peer Computational Grid Framework. Lambert Academic Publishing (2010)
4. Bent, J., Thain, D., Arpaci-Dusseau, A.C., Arpaci-Dusseau, R.H., Livny, M.: Explicit control in a batch-aware distributed file system (March 2004)
5. Callback File System, <http://www.eldos.com/cbfs/>
6. Condor project, <http://www.cs.wisc.edu/condor/>
7. D4Science: Distributed collaboratories Infrastructure on Grid ENabled Technology 4 Science, <http://www.d4science.eu/>
8. Dokan file system, <http://dokan-dev.net/en/>
9. EGEE: Enabling Grids for E-Science in Europe, <http://public.eu-egee.org/>
10. Ernst, M., Fuhrmann, P., Gasthuber, M., Mkrtchyan, T., dCache, C.W.: A distributed storage data caching system. In: Computing in High Energy Physics, Beijing, China (2001)
11. Filesystem in Userspace, <http://fuse.sourceforge.net/>

12. GLite: Light weight middleware for grid computing, <http://glite.web.cern.ch/glite/>
13. The Globus toolkit, <http://www.globus.org/toolkit/>
14. GridFTP, <http://globus.org/toolkit/data/gridftp/>
15. Jones, M.B.: Interposition agents: Transparently interposing user code at the system interface. In: fourteenth ACM symposium on Operating systems principles, Asheville, North Carolina, United States, pp. 80–93 (1994)
16. Lederer, H., Pringle, G.J., Girou, D., Hermanns, M.-A., Erbacci, G.: Deisa: Extreme computing in an advanced supercomputing environment (2007)
17. Leech, M., Ganis, M., Lee, Y., Kuris, R., Koblas, D., Jones, L.: Socks protocol version 5 (1996)
18. MacFUSE, <http://code.google.com/p/macfuse/>
19. Miller, B., Callaghan, M., Cargille, J., Hollingsworth, J., Irvin, R.B., Karavanic, K., Kunchithapadam, K., Newhall, T.: The paradyn parallel performance measurement tools. *IEEE Computer* 28(11), 37–46 (1995)
20. Narasimhan, P., Moser, L.E., Melliar-Smith, P.M.: Exploiting the internet inter-orb protocol interface to provide corba with fault tolerance. In: *Unix Conference on Object-Oriented Technologies and Systems*, Portland, Oregon (1997)
21. NorduGrid: Nordic Testbed for Wide Area Computing and Data Handling, <http://www.nordugrid.org/>
22. Pike, R., Presotto, D., Thompson, K., Trickey, H., Winterbottom, P.: The use of name spaces in Plan 9. *SIGOPS Oper. Syst. Rev.* 27(2), 72–76 (1993)
23. G.R.: Grid3: An application grid laboratory for science. In: *Computing in High Energy Physics and Nuclear Physics*, p. 18. Interlaken, Switzerland (2004)
24. Son, S., Livny, M.: Recovering internet symmetry in distributed computing. In: *CC-Grid*, Los Alamitos, CA, IEEE Computer Society, Los Alamitos (2003)
25. Thain, D., Livny, M.: Parrot: Transparent user-level middleware for data-intensive computing. *Scalable Computing: Practice and Experience* 6(3), 9–18 (2005)
26. UNICORE: Uniform Interface to Computing Resources, <http://www.unicore.eu>
27. van Hensbergen, E.V., Evans, N.P., Stanley-Marbell, P.: A unified execution model for cloud computing. In: *Large Scale Distributed Systems and Middleware (LADIS 2009)*, October 2009, ACM, New York (2009)