

The Globus eXtensible Input/Output System (XIO): A protocol independent IO system for the Grid

William Allcock, John Bresnahan, Rajkumar Kettimuthu and Joseph Link
Mathematics and Computer Science Division
Argonne National Laboratory
Argonne, Illinois 60439
Email: {allcock,bresnaha,kettimut,link}@mcs.anl.gov

Abstract— In distributed heterogeneous Grid environments the protocols used to exchange bits are crucial. As researchers work hard to discover the best new protocol for the Grid, application developers struggle with ways to use these new protocols. A stable, consistent, and intuitive framework is needed to aid in the implementation and use of these protocols. While the application must not be burdened with the protocol details some of it may need to be exposed to take advantage of potential optimizations. In this paper we examine how the Globus XIO API provides this framework. We will explore the performance implications of using this abstraction layer and the benefits gained in application as well as protocol development.

I. INTRODUCTION

It is amazing – or frustrating, depending on your point of view – how many I/O problems appear to be solvable with simple Open/Close/Read/Write (OCRW) functionality. The devil is in the details, however, and the solutions to these problems are just different enough to result in a variety of API models with varying sets of semantics. From a very broad perspective they all look alike, since they all have a source and a sink between which a stream of data flows. However, a more detailed examination shows differences that are extremely significant to the application developer. For example, if the end goal of an application is to transfer files between two locations, HTTP [1] and GridFTP [2]–[5] may be logical choices. The choice between them may be based on performance and environment characteristics, but the architecture and design of the application is the same. However when looked at from the perspective of a developer actually implementing that design, the asynchronous push model FTP API, and the synchronous multi-threaded API for HTTP messaging is entirely different and drastically affects the design of the application.

The Grid [6], [7] is a heterogeneous, dynamic and evolving environment. Development of communication protocols to optimize use of and access to the underlying fabric of the Grid is a key area of research. This research turns out protocols that have similar goals but with varying properties ideal for different environments. As researchers develop new transfer and control protocols, application developers struggle

to update their applications in order to gain the benefits of increased network utilization. In order to make their applications portable to many heterogeneous grid environments, the developers may try to integrate as many protocols as possible. This typically means many different APIs all with different semantics, many different programming models, and even different sets of bugs to work around. This issue certainly leads to a longer development cycle and inevitably leads to messy code.

On the other end of this problem is the creation of new protocols. As researchers and developers create new protocols the focus of their effort should be on the details of their protocol, yet too often a lot of time is spent on side issues. The development of these new protocols can be significantly delayed by issues that have more to do with user interface than they do with the protocol itself. Designing a good intuitive API is a task in itself, and one that is unfortunately often neglected. Further, the majority of bugs and development time in a robust full featured protocol library are results of race conditions due to multi-threading and event notification. These issues are paramount to the implementation of a protocol but not central to the protocol itself. Additionally, most new protocols are layered on top of standard transport protocols, like TCP [8]. This leaves the developer coding directly to the socket library, which is a difficult art to master. Globus XIO attempts to provide a framework that aids in the implementation and use of the evolving protocols. It helps to minimize the time and effort involved in implementing new protocols. Applications that are written using Globus XIO can make use of these new protocols with very little or no changes to the application code.

The rest of this paper is organized as follows. Section II provides an overview of Globus XIO. We describe a common usage scenario in Section III and the architecture of XIO in Section IV. Section V and Section VI explain the user API and the UDT (UDP based Data Transfer protocol) driver written using XIO framework respectively. We analyze the performance of Globus XIO in Section VII. After discussing related work in Section VIII, we summarize in Section IX.

II. OVERVIEW OF GLOBUS XIO

Globus XIO is a simple OCRW abstraction layer to transport protocols that was designed with these problems in mind. To application developers, Globus XIO is an API that enables different I/O problems to be presented uniformly as a simple OCRW interface with a single set of semantics. To network protocol developers, it is a support framework for developing communication protocols. To mass storage vendors, it is an interface that enables an existing application written with XIO to access their hardware. Globus XIO does not try to solve every I/O problem by mapping every type of I/O to a single API; that is clearly not possible. Nor does Globus XIO attempt to hide all details of the underlying protocol; that is clearly undesirable. While it is nice to think that all I/O can be lumped into a single abstraction, protocols are different for a reason. They solve different problems and do different jobs. Often a user must be able to control these differences to properly take full advantage of the I/O protocol involved.

The first problem that Globus XIO addresses is the abstraction of byte stream IO into a simple OCRW API. It is unnatural (and unwise) to attempt to map all I/O types to an OCRW API. However, those that can be viewed as streams of bytes, which is a significant number of them, map very well. All the differences cannot be hidden, but the impact on the application code can be minimized and the differences isolated. So whether data is read from a file or a network or an electron microscope, whether it is using TCP, UDP [9], or some custom protocol, to the application using Globus XIO the interface is simply OCRW. Globus XIO also helps to minimize the development time and effort involved in creating and prototyping new I/O protocols and new device interfaces. Globus XIO provides a framework that handles many of the non-protocol specific ancillary requirements of a protocol developer such as event handling, error handling, etc. The protocol specific functionality is abstracted into a driver, where the protocol developer can focus on the interesting details of the protocol they are prototyping or producing for production. The driver must implement a well-defined set of interface functions, with a well-defined set of semantics associated with them. Essentially, these functions consist of Open, Close, Read, and Write. For instance, a driver developer defines an open function, and implements the code specific to their protocol for opening. The same steps are taken for the read, write and close functions. When implementing these functions the developer does not need to be concerned with user error checking, this has all been done by the framework and all parameters are guaranteed to be good. This and all other user interface issues are removed from the concern of the driver developer.

Often, APIs are created according to a blocking model. This has the advantage of being much easier to design and imple-

ment, however it will quite possibly have performance issues. It is likely that while an application is waiting for bytes to arrive via some transport that it can perform data processing. If using a blocking API the only way an application can do these parallel tasks is to have separate threads for IO and processing. This can easily lead to complicated synchronization issues that a scientific application may not be willing to deal with. Further, it does not scale well as more and more simultaneous IO occurs. To overcome this issue the protocol developer must create their own polling loop, which can be very complicated low level code and is nothing more than a distraction to the protocol developer. The Globus XIO framework alleviates this problem from the driver developer by providing asynchronous functionality. A rich set of system polling code is distributed with XIO, and an internal API provides the driver with polling and barriers and other asynchronous functionality.

One important feature of XIO is that drivers can also use other drivers. This means that once a TCP driver is created (and one is distributed with Globus XIO) no other driver needs to code directly to sockets again, they can instead use the TCP driver. So when creating the HTTP driver, the details of framing HTTP messages can be focused on and the work of transporting the messages can simply be passed onto the TCP driver.

III. A COMMON USAGE SCENARIO

One of the common Grid computing usage scenarios is referred to as a “DataGrid” problem [2]. This involves accessing a remote data source, transporting it to a compute node, performing computation on it, and then sending the results to a remote destination. This scenario involves multiple IO operations, each of which may involve a different programming model, API, etc. We will examine this scenario in some detail to see how XIO might be applied.

In our scenario, a remote data source is accessed for input data to the computation. Common access methodologies for this data might include GridFTP, HTTP, or a proprietary mass storage system such as HPSS [10], [11]. However, the source of data could also be something more unusual such as a scientific instrument of some kind or even the output of another computational program. In all these cases, this IO can be easily represented as a byte stream, yet the programming model for each could (and likely would) be quite different. If the application wishes to be able to access all these sources, the application must contain code for all these proprietary APIs and deal with all of the different programming models. However, if we interpose XIO as an abstraction layer, the application sees a uniform programming model, uniform API, uniform error semantics, etc. Obviously, the differences are still present, but they have been isolated to the driver, making

the application logic much cleaner and localizing the location of source specific issues.

Our scenario also calls for sending the results of the computation to a remote destination. It might be the same or similar access methods as already described, but what if it were a visualization display of some kind? This case is much less clear and it depends on the characteristics of the display. If the display were some sort of movie viewer and expected a simple video stream, then XIO might well be appropriate. However, if complex windowing operations were required, XIO would definitely NOT be appropriate.

IV. BASIC ARCHITECTURE

There are two important concepts in the architecture of Globus XIO, the driver and the stack. We discussed the driver above; it is the abstracted component where all the protocol specific implementation exists. There are two types of drivers, transform drivers and transport drivers. Transport drivers are those that actually move data into or out of the process space. Examples of this are TCP and UDP. Transform drivers are those that manipulate, examine, frame, or change the data, or in other words, drivers that take any action other than moving the data across the process boundary. Examples would be compression, logging, and HTTP.

A particular protocol implementation can involve many drivers; this is where the concept of a stack of drivers comes into play. In a stack there must only be one transport driver, and it must be at the bottom of the stack. This stands to reason since the transport driver is what actually moves the bits on a wire. There can be any number of transform drivers on a driver stack. As data is written it moves from the user to the first driver on the stack. The data is moved down the stack, through the effects of each transform driver (if any are used) to the transport driver. The transport driver then ships the data over the wire. Data follows a similar path for reads, only in the opposite direction, from wire to transport driver, and up the stack to the user. The protocols that make use of multiple transport drivers have to construct multiple stacks. For example, a protocol that uses TCP for exchanging control information and UDP for transferring the actual data, has to have two different stacks: one with TCP as a transport and the other with UDP as the transport driver.

Driver stacks can be mixed and matched. An HTTP driver has been created as a transform driver, and typically it sits in a stack directly above TCP. For example, if a driver is created implementing UDT (a reliability layer over UDP), a stack can be formed of HTTP on top of UDT without a single code change to any of the drivers involved. A protocol may provide some special features that other protocols may not support. Globus XIO provides a way for the protocol developers to expose the functionalities that are specific to their protocols.

Each driver may have its own attribute structure. Globus XIO provides user API to control the attributes of a driver. It gives the user an opportunity to tweak parameters that are specific to a driver. The attribute support is optional and a driver may choose not to have attribute support at all.

In Globus XIO there are no extra memory copies. The user passes in a buffer to the framework. The framework then passes a pointer to that buffer down the driver stack. If any driver along the way needs to insert a header or append a footer it can make use of the `readv/writev` functions that XIO provides. If a driver needs to alter the data (as a compression driver would) then a copy is required, however this is a mandatory copy to achieve the desired functionality, it is not an extra copy.

V. THE GLOBUS XIO USER API

Our discussion up to this point has been fairly abstract. At this point we will give a brief overview of the status of Globus XIO and some code snippets demonstrating the use of the user API. Globus XIO is available in the latest stable version of the Globus Toolkit (Globus Toolkit 3.2). Documentation for the entire API is available at [12]. Globus XIO is a C library. Since the core API is simply OCRW, the user API is fairly straightforward. We provide synchronous calls (`globus_xio_read()`), asynchronous calls (`globus_xio_register_read()`) and have vector variants of each (`globus_xio_[register]_readv()`). The best way to become familiar with Globus XIO is by looking at some code snippets.

Step 1: Activate Globus

XIO was developed by the Globus Alliance and so inherits some Globus Toolkit semantics. Accordingly, as with all Globus Toolkit programs, you must first activate the Globus module. Until activation is complete, no XIO function calls can be successfully executed. The module is activated with the following line:

```
globus_module_activate(GLOBUS_XIO_MODULE);
```

Step 2: Load Driver

The next step is to load all the drivers needed to complete the I/O operations in which you are interested. The function `globus_xio_load_driver()` is used to load a driver. To successfully call this function, you must know the name of all the drivers you wish to load. For this example we want to load only the file io driver. The prepackaged file io driver's name is "file." This driver is loaded with the following code:

```
globus_result_t res;  
globus_xio_driver_t driver;  
res = globus_xio_load_driver(&driver, "file");
```

If upon completion of the above function call, the variable `res` is equal to `GLOBUS_SUCCESS`, then the driver was

successfully loaded and can be referenced with the variable “driver.”

Step 3: Create Stack

Once `globus_xio` is activated and a driver loaded, you need to build a driver stack. In our example the stack consists of only one driver, the file driver. The stack is established with the following code (building from the above code snips):

```
globus_xio_stack_t stack;
globus_xio_stack_init(&stack);
globus_xio_stack_push_driver(stack, driver);
```

Step 4: Opening the Handle

Once the stack is created, you can open a handle to the file, in one of two ways. The first way is a passive open. An example is a TCP listener. The open is performed passively and then it waits for some other entity to act on it. The other alternative is an active open. An example is a TCP connect. The user, who initiates the open, performs the open actively. Our example has an active open. To create a handle for an active open, you first initialize the handle object and then open it with the contact information. The following code illustrates this:

```
globus_xio_handle_t handle;
globus_xio_handle_create(&handle, stack);
res = globus_xio_open(handle, "/tmp/junk.txt", NULL);
```

Step 5: The Payoff

Now that you have an open handle to a file, you can read or write data to it with either: `globus_xio_read()` or `globus_xio_write()`. Once you are finished performing I/O operations on the handle, you should call `globus_xio_close(handle)`.

All this may seem like a lot of effort for simply reading a file. The advantages become clear, however, when you wish to use other drivers. In the above example, it would be trivial to change the I/O operations from file I/O to TCP, HTTP, or FTP. All you would need to do is change the driver name string passed to `globus_xio_load_driver()` and change the contact string passed to `globus_xio_[register]_open()`. Both can be done easily at runtime, as the program `globus_xio_example.c` in the appendix demonstrates.

VI. THE UDT DRIVER

In network protocol research a common goal is optimal bandwidth utilization, while still being network friendly. The drawback of TCP inherent in its AIMD based congestion control mechanism [13]–[15] is well known [16]. Furthermore, if the bandwidth-delay product is very high, it takes a very long time for TCP to recover from a packet loss. Researchers have come up with numerous alternatives. These solutions include improvements to TCP [17]–[20], new transport protocols such

as XCP [21], XTP [22] and reliable layers on top of UDP [23]–[27]. Globus XIO is the perfect framework for researching reliability layers on top of UDP. A UDP driver already exists, so all of the socket code is taken care of and the framework assists greatly in quick prototyping of protocols. Further, since the user API does not change based on the driver used, a single set of performance measuring tools can be used to compare many drivers.

One particular protocol in this category is UDT (developed at the Laboratory for Advanced Computing at University of Illinois at Chicago) [26]. A UDT driver was implemented for the Globus XIO framework. Globus XIO allows the driver developer to load and control the desired protocol underneath the driver. This feature is very useful for developing drivers that are in an evolving state like SABUL/UDT. SABUL [25] was the first iteration of UDT and used TCP for exchanging control messages and UDP for transferring the data. Later, an improved version of SABUL, namely UDT, which use UDP for transferring both control and data, was proposed. These changes were made easily in Globus XIO by simply altering the driver loaded. This would have been a much more substantial effort if the development were done using the socket library.

The provision of both synchronous and asynchronous API and the framework handling the technicalities of that made the development of UDT easier. There was little worry of synchronizing processing by waiting and signaling. Instead the framework abstracted user requests for data into operations. When a request to read data from UDT was made it was given an operation. UDT then checks its read buffer and if it had already read the requested amount of data, it finishes the operation immediately. Otherwise it puts the operation in the read operation queue. Similarly, when a request to write is made, the buffer and the operation would be put in the queue. It then will split the data into multiple packets, add header and send it to the destination. Whenever an operation was complete, the driver simply notified the framework and the framework took care of all functionality necessary to safely notify the user. Additionally, the Globus XIO framework allows the drivers to initiate new asynchronous operations that are performed in parallel to user initiated operations. This is an important feature for UDT, and any UDP based reliable transport protocols, as they need to allow for the protocol layer to continuously read the data. Without this feature of Globus XIO the protocol developer would have to use synchronous operations and manage threads internally. Also, when the user initiates a close, the XIO framework makes sure all the outstanding operations complete/cancel and its callback called before the close callback is called. This makes the driver developer’s job easier by relieving him from worrying about canceling or waiting for the outstanding operations while

performing the close operation.

Driver Reuse

To illustrate the reuse of drivers let us look at a usage scenario with the UDT driver described above. The UDT driver provides the same semantics as TCP to the user of XIO. Given those same semantics the user may wish to have the same security mechanisms. In other words, if it is desirable to have a GSI authenticated TCP connection, it is also desirable to have a GSI authenticated UDT connection. With the XIO framework, reuse of the security code is trivial. To have a secure TCP channel the user simply needs to build a stack with the GSI driver on top of the TCP driver. Similarly, with UDT, a stack is built of the GSI driver on top of the UDT driver. There are no alterations to the user’s code, and no alterations to the driver code. The needed drivers are simply plugged into the framework.

VII. PERFORMANCE RESULTS

In computer science there is the old adage that there is no problem you cannot solve by adding a layer of abstraction, and there is no optimization you cannot make by removing one. Since Globus XIO is an abstraction layer to protocols it is important to look at its performance characteristics. We have conducted several performance tests bench marking Globus XIO against Globus IO (an API with similar features, yet it has been coded directly against the socket library), and the standard socket library. We did both latency and bandwidth tests on a local host, between two different hosts on a LAN, between two hosts connected over a wide area link.

A. Latency Tests

The latency tests were carried out in a ping-pong fashion. The sender sends a message with a certain data size to the receiver and waits for a reply from the receiver. The receiver receives the message from the sender and sends back a reply with the same data size. Each ping-pong test was carried out 1000 times, the total time to complete those iterations was then used to determine an average one-way latency number. Figures 1 through 3 show the results of the latency tests. For tests conducted on the loopback interface, shown in Fig. 1, the latency of a single byte increased from 20 us for raw socket IO to 37 us for IO via Globus XIO. However, the percentage increase in latency decreases as the message size increases. For a message of size of 100KB, the percentage increase is 8% and for a message size of 10 MB, it is just 2%. Tests conducted on the LAN provide much better results, as shown in Fig. 2. The percentage increase in latency for Globus XIO is 8% for transferring a 1000 byte message and it decreases as the message size increases and it is only 0.003% for a message of size of 10MB. Finally, Fig. 3 shows the latency numbers

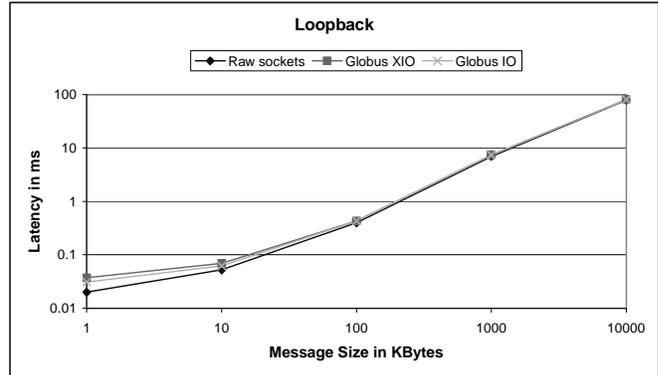


Fig. 1. Comparison of one-way latency for transferring messages of various sizes over the loopback interface on a local host.

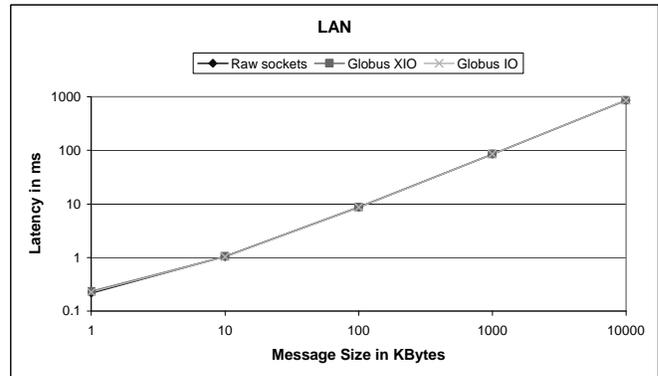


Fig. 2. Comparison of one-way latency for transferring messages of various sizes between two hosts on a LAN.

over a WAN. It is here that Globus XIO comes into its own. The overhead introduced is less than 1.5% for any message size.

In general, for most Globus applications, latency is not a critical issue and the use of Globus XIO will be of negligible impact. In certain latency critical applications, such as MPI application the utility of XIO would need to be determined with more exhaustive tests.

B. Bandwidth Tests

The bandwidth tests were carried out by having the sender sending out a large number (1000) of back-to-back messages to the receiver and then waiting for a reply from the receiver. The receiver sends the reply only after receiving all (1000) messages. Then the bandwidth was calculated based on the elapsed time (from the time sender sends the first message until the time it receives the reply back from the receiver) and the number of bytes sent by the sender. Figures 4 and 5 indicate that the throughput achieved with Globus XIO is nearly the same as that of raw sockets. The percentage decrease in the throughput is less than 3% in all cases, and for large messages it is less than 1.5%. While some applications may not be able

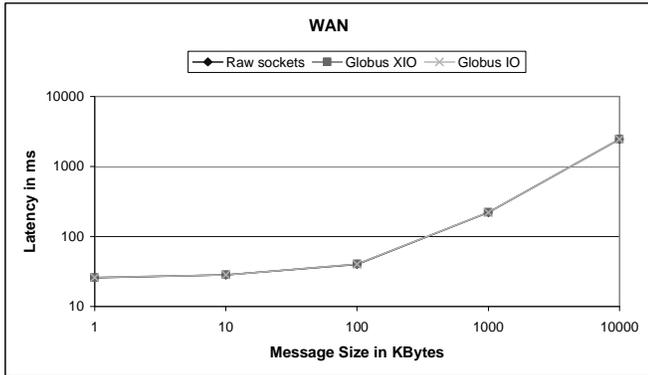


Fig. 3. Comparison of one-way latency for transferring messages of various sizes between two hosts connected over a wide area link.

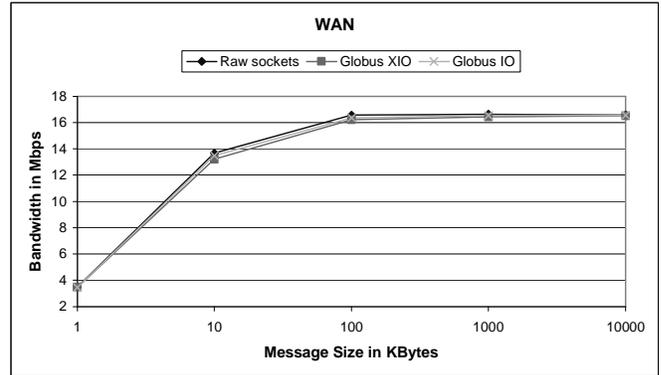


Fig. 5. Comparison of throughput achieved over a wide area network where the bottleneck is an OC-12 link.

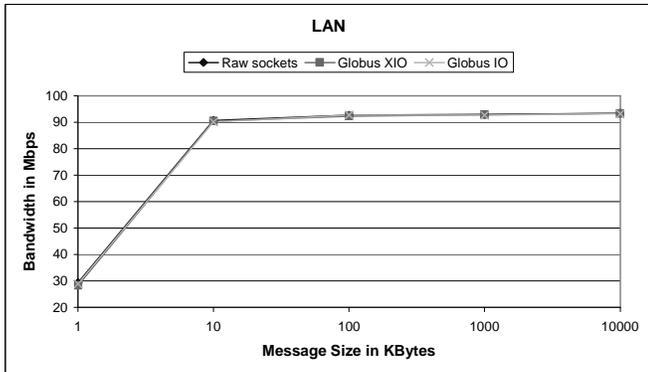


Fig. 4. Comparison of throughput achieved over a gigabit link on a local area network.

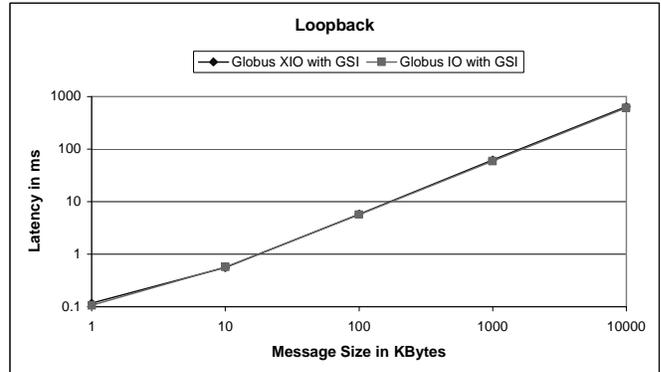


Fig. 6. Comparison of latency for Globus XIO over GSI with Globus IO over GSI over the loopback interface on a local host.

to tolerate this loss in bandwidth, in general, the losses are negligible in a real world application.

All the tests shown in the figures above were conducted with a single XIO driver loaded. However, XIO was designed to have a “stack” of multiple drivers and in some cases this is required to obtain equivalent functionality. Figure 6 shows a comparison of Globus IO using GSI authentication against Globus XIO using a stack consisting of a TCP driver and a GSI driver. The latency overhead is still minimal at less than 5% for small messages and less than 2% for larger messages.

VIII. RELATED WORK

Many operating systems support abstractions for encapsulating protocols; examples include Berkeley Unix sockets [28] and System V Unix streams [29]. Such abstractions are useful because they provide a common interface to a collection of dissimilar protocols available in the kernel space, thereby simplifying the task of composing protocols. In a grid environment, researchers have come up with lots of application level protocols and there is no standard abstraction to compose all these protocols. Globus XIO provides this abstraction and the application developers can use this framework to easily switch

between the protocols. X-kernel [30] provides an object-based framework for implementing and testing new protocols. It is actually a new operating system kernel that provides an explicit architecture for constructing and composing protocols whereas Globus XIO is a framework that eases the job of the application protocol developers by providing a simple interface to the underlying protocols. The Pluggable file system (now renamed as parrot virtual file system) [31] transforms an application’s standard POSIX I/O operations into a variety of remote I/O operations. It allows the user to select the desired protocol and perform I/O operations on local or remote file system. It also provides an interactive interface to the user. Though it has some features that are present in Globus XIO, it does not provide any interface for developing new drivers. Also, it does not allow the users to build a stack with multiple protocols.

IX. SUMMARY

Globus XIO provides a framework that implements a simple Open/Close/Read/Write interface that is appropriate for most byte stream oriented IO applications. The framework user API presents a uniform API, semantics, and error handling to

the application. Variations in programming models due to the different underlying protocols and APIs are isolated in drivers and drivers may be “stacked” to get composite functionality. Our early performance tests show that the overhead associated with this abstraction and the resulting benefits is minimal, particularly in most real world situations, generally less than 3%.

The primary benefits to be gained from using Globus XIO include:

A uniform IO API that developers can learn and be comfortable with:

A single, rational set of APIs, semantics, error responses, as well as problems and work arounds. Differences due to different network protocols or data access methods are isolated in the drivers making the application more stable and minimizing the programmer involvement with these differences.

Ability to quickly adapt to different transport protocols:

Applications written with XIO can switch from TCP to other protocols with minimal changes to the application.

Ability to quickly adapt to different data access mechanisms:

Files stored behind GridFTP servers, on HPSS or UniTree mass storage devices or any number of other vendor proprietary storage systems can all be accessed in uniform manner. This also provides providers of these access mechanisms a convenient vehicle for providing this access and immediately makes their access mechanisms available to any application utilizing XIO.

Provides a convenient framework for research and development:

Researchers in new protocols and data access mechanisms can utilize the XIO framework during their research and development cycles. The XIO framework provides many of the facilities required to successfully conduct their research, but outside their interest and expertise.

APPENDIX

GLOBUS_XIO_EXAMPLE.C

```
#include "globus_xio.h"

int
main(
    int      argc,
    char *   argv[])
{
    globus_result_t  res;
    char *           driver_name;
```

```
globus_xio_driver_t  driver;
globus_xio_stack_t   stack;
globus_xio_handle_t  handle;
globus_size_t        nbytes;
char *               contact_string = NULL;
char                 buf[256];

contact_string = argv[1];
driver_name = argv[2];

globus_module_activate(GLOBUS_XIO_MODULE);
res = globus_xio_driver_load(
    driver_name,
    &driver);
assert(res == GLOBUS_SUCCESS);
res = globus_xio_stack_init(&stack, NULL);
assert(res == GLOBUS_SUCCESS);
res = globus_xio_stack_push_driver(stack, driver);
assert(res == GLOBUS_SUCCESS);
res = globus_xio_handle_create(&handle, stack);
assert(res == GLOBUS_SUCCESS);
res = globus_xio_open(handle, contact_string, NULL);
assert(res == GLOBUS_SUCCESS);
do
{
    res = globus_xio_read(handle, buf, sizeof(buf) - 1, 1,
        &nbytes, NULL);
    if(nbytes > 0)
    {
        buf[nbytes] = '\0';
        fprintf(stderr, "%s", buf);
    }
} while(res == GLOBUS_SUCCESS);
globus_xio_close(handle, NULL);
globus_module_deactivate(GLOBUS_XIO_MODULE);
return 0;
}
```

ACKNOWLEDGMENTS

This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy, under Contract W-31-109-ENG-38.

REFERENCES

- [1] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Manister, P. Leach, and T. Berners-Lee, “Hypertext Transfer Protocol - HTTP/1.1. IETF, RFC 2616,” June.
- [2] W. Allcock, J. Bester, J. Bresnahan, A. Chervenak, I. Foster, C. Kesselman, S. Meder, V. Nefedova, D. Quesnel, and S. Tuecke, “Data management and transfer in high performance computational grid environments,” *Parallel Computing Journal*, vol. 28, no. 5, pp. 749–771, 2002. [Online]. Available: <http://www.globus.org/research/papers/dataMgmt.pdf>

- [3] —, “Secure, efficient data transport and replica management for high-performance data-intensive computing,” in *Proceedings of the 18th IEEE Symposium on Mass Storage Systems, San Diego, California, April 17-20, 2001*, pp. 13–28. [Online]. Available: <http://www.globus.org/research/papers/msc01.pdf>
- [4] W. Allcock, “GridFTP: Protocol extensions to FTP for the grid. Global Grid ForumGFD-R-P.020, 2003.” [Online]. Available: <http://www.ggf.org/documents/GWD-R/GFD-R.020.pdf>
- [5] W. Allcock, A. Chervenak, I. Foster, L. Pearlman, V. Welch, and M. Wilde, “Globus toolkit support for distributed data-intensive science,” in *Proceedings of International Conference on Computing in High Energy and Nuclear Physics, Beijing, China, September 2001*. [Online]. Available: <http://www.globus.org/research/papers/Globus.CHEP01.pdf>
- [6] I. Foster, C. Kesselman, and S. Tuecke, “The anatomy of the Grid: Enabling scalable virtual organization,” *The International Journal of High Performance Computing Applications*, vol. 15, no. 3, pp. 200–222, Fall 2001.
- [7] I. Foster and C. Kesselman, *The grid: blueprint for a new computing infrastructure*. Morgan Kaufmann Publishers Inc., 1999.
- [8] J. Postel, “RFC 793: Transmission Control Protocol,” September 1981.
- [9] —, “RFC 768: User Datagram Protocol,” August 1980.
- [10] “High performance storage system.” [Online]. Available: <http://www.hpss-collaboration.org/hpss/>
- [11] D. Teaff, D. Watson, and B. Coyne, “The architecture of the high performance storage system (hpss),” 1995. [Online]. Available: citeseer.ist.psu.edu/teaff95architecture.html
- [12] “Globus XIO documentation.” [Online]. Available: <http://www-unix.globus.org/developer/xio/>
- [13] V. Jacobson, “Congestion Avoidance and Control,” *ACM SIGCOMM Computer Communication Review*, vol. 18, no. 4, pp. 314–329, 1988.
- [14] M. Allman, V. Paxson, and W. Stevens, “RFC 2581: TCP Congestion Control,” 1999.
- [15] S. Floyd and T. Henderson, “RFC 2582: The NewReno Modification to TCP’s Fast Recovery Algorithm,” 1999.
- [16] D. Katabi, M. Handley, and C. Rohrs, “Internet congestion control for future high bandwidth-delay product environments,” 2002. [Online]. Available: citeseer.ist.psu.edu/katabi02internet.html
- [17] S. Floyd, “HighSpeed TCP for large congestion windows RFC 3649, Experimental December 2003.” [Online]. Available: <http://www.faqs.org/rfcs/rfc3649.html>
- [18] T. Kelly, “Scalable TCP: improving performance in highspeed wide area networks , journal = SIGCOMM Computer Communication Review, volume = 33, number = 2, year = 2003, pages = 83–91, doi = <http://doi.acm.org/10.1145/956981.956989>, publisher = ACM Press,.”
- [19] C. Jin, D. X. Wei, and S. H. Low, “FAST TCP: motivation, architecture, algorithms, performance,” in *Proceedings of the IEEE Infocom*, March 2004.
- [20] D. J. Leith and R. Shorten, “H-TCP protocol for high-speed long distance networks,” in *Second International Workshop on Protocols for Fast Long-Distance Networks, Argonne, IL, February 16-17, 2004*.
- [21] D. Katabi, M. Handley, and C. Rohrs, “Congestion control for high bandwidth-delay product networks,” in *Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*. ACM Press, 2002, pp. 89–102.
- [22] W. T. Strayer, M. Lewis, and R. E. Cline, Jr., “XTP as a transport protocol for distributed parallel processing,” in *Proceedings of the USENIX Symposium on High-Speed Networking, oakland, CA, August 1994*, pp. 91–101. [Online]. Available: citeseer.nj.nec.com/strayer94xtp.html
- [23] D. Clark, M. Lambert, and L. Zhang, “NETBLT: A Bulk Data Transfer Protocol. IETF, RFC 998,” March 1987.
- [24] E. He, J. Leigh, O. Yu, and T. A. DeFanti, “Reliable Blast UDP : Predictable high performance bulk data transfer,” in *Proceedings of the IEEE International Conference on Cluster Computing (CLUSTER’02), Chicago, IL, September 23 - 26, 2002*, pp. 317–324. [Online]. Available: <http://www.evl.uic.edu/cavern/papers/cluster2002.pdf>
- [25] H. Sivakumar, R. L. Grossman, M. Mazzucco, Y. Pan, and Q. Zhang, “Simple available bandwidth utilization library for high-speed wide area networks,” *Journal of Supercomputing*, 2004.
- [26] Y. Gu and R. Grossman, “UDT (UDP based Data Transfer Protocol): An Application Level Transport Protocol for Grid Computing,” in *Second International Workshop on Protocols for Fast Long-Distance Networks, Argonne, IL, 2004*.
- [27] “Tsunami network protocol implementation.” [Online]. Available: <http://www.indiana.edu/~anml/anmlresearch.html>
- [28] S. J. Leffler, W. N. Joy, R. S. Fabry, and M. J. Karels, “Networking implementation notes 4.4BSD edition.” [Online]. Available: citeseer.nj.nec.com/28738.html
- [29] C. AT&T, *UNIX system V: streams programmer’s guide*. Prentice-Hall, Inc., 1987.
- [30] N. C. Hutchinson and L. L. Peterson, “The x-kernel: An architecture for implementing network protocols,” *IEEE Transactions on Software Engineering*, vol. 17, no. 1, pp. 64–76, 1991. [Online]. Available: <http://www.cs.arizona.edu/xkernel/bibliography.html#Hutchinson91tse>
- [31] “Pluggable file system.” [Online]. Available: <http://www.cs.wisc.edu/condor/pfs/>