

Community Software Development with the Astrophysics Simulation Collaboratory

⁵Gregor von Laszewski, ²Michael Russell, ^{2,5}Ian Foster, ^{3,4}John Shalf, ¹Gabrielle Allen, ⁴Greg Daues, ³Jason Novotny, ^{1,4}Edward Seidel

1 Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut, Golm, Germany

2 University of Chicago, Chicago, IL, U.S.A.

3 Lawrence Berkeley National Laboratory, Berkeley, CA, U.S.A.

4 National Center for Supercomputing Applications, Champaign, IL, U.S.A.

5 Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL, U.S.A.

6 Washington University, St Louis, MI, U.S.A.

Corresponding author: Gregor von Laszewski: gregor@mcs.anl.gov, phone 630 252 0472, fax 630 252 5896

Submitted to:

*Journals Production Department, John Wiley & Sons, Ltd.,
Chichester, West Sussex, PO19 1UD, U.K.*

SUMMARY

We describe a Grid-based collaboratory that supports the collaborative development and use of advanced simulation codes. Our implementation of this collaboratory uses a mix of Web technologies (for thin-client access) and Grid services (for secure remote access to, and management of, distributed resources). Our collaboratory enables researchers in geographically disperse locations to share and access compute, storage, and code resources, without regard to institutional boundaries. Specialized services support community code development, via specialized Grid services, such as online code repositories. We use this framework to construct the Astrophysics Simulation Collaboratory, a domain-specific collaboratory for the astrophysics simulation community. This Grid-based collaboratory enables researchers in the field of numerical relativity to study astrophysical phenomena by using the Cactus computational toolkit.

KEY WORDS: Collaboratory, Grid Computing, Globus, Cactus, Java CoG Kit, Astrophysics Simulation Collaboratory

1 INTRODUCTION

Solving large complex problems often requires coordinated effort by geographically distributed participants. The term *collaboratory* was created in the early 1990s to

denote an infrastructure supporting such collaborative endeavors [19]. Typically, collaboratories have emphasized interpersonal communication and collaborative access to relatively simple resources. We are interested in extending the collaboratory concept to address problem-solving methods and environments that require the sharing and coordinated use of large numbers of distributed resources. In this context, we see considerable value for *Grid* technologies, which address authentication, resource discovery, and resource access issues that arise when sharing distributed resources. Grids have emerged over the past decade as a distributed computing infrastructure to promote the sharing of hardware, information, knowledge, and expertise across institutional and disciplinary boundaries [13, 14]. We term this combination of collaborative and Grid technologies a *Grid-based collaboratory*.

Although the results of our research can be applied to a variety of domains, we focus our attention here on computational astrophysics, a discipline that has challenging requirements for large-scale collaborative development and application of simulation, data analysis tools. For example, a typical simulation experiment might involve the development of a simulation code from a dozen distinct components, developed by different groups; the acquisition of computer resources at one or more remote sites; monitoring and steering of the computation as it runs; and access to data resources at one or more output locations. Our goal is to develop an *astrophysics simulation collaboratory* (ASC) that allows astrophysics and computational science researchers to cooperate as they develop and share software, computational, and data resources.

Discipline scientists are typically not experts in distributed computing. Hence, it is necessary that our collaboratory hide, as much as possible, the complexity of using multiple remote resources. In the past decade Web technologies have been accepted by many scientific communities as a means of distributing and accessing services. Web pages present the information needed in an orderly and structured fashion; Web protocols can be used to access remote information and services. Yet at the same time, Web protocols and technologies are inadequate for the sharing and coordination issues that arise when running simulations on remote systems, sharing multiple resources, and so forth. Here, Grid technologies [14] can play an important role. Commodity Grid (CoG) Kits [22] provide components for accessing Grid services with commodity technologies such as Java, JSP, CORBA, and Python.

In this paper we describe both the framework, services, and components needed to construct Grid-enabled collaboratories, and the particular application of these technologies within our ASC. The ASC [7, 21] fulfills the following general requirements posed by the astrophysics simulation community:

- **Promote the creation of a community** for sharing and developing simulation codes and scientific results. The Grid infrastructure promotes the creation of virtual organizations [14] (VOs) within which resources are shared across institutional boundaries through Grid services.
- **Enable transparent access to remote** resources, including computers, data storage archives, information servers, and shared code repositories.
- **Enhance domain-specific component and service development** supporting specific problem-solving capabilities, such as the development of simulation codes for the astrophysical community or the development of advanced Grid services reusable by the community.
- **Distribute and install programs onto remote resources** while accessing code repositories, compilation, and deployment services.
- **Enable collaboration during program execution** to foster interaction during the development of parameters and the verification of the simulations.
- **Enable shared control and steering of the simulations** to support asynchronous collaborative techniques among collaboratory members.
- **Provide access to domain-specific clients** that, for example, enable access to multimedia streams and other data generated during the execution of the simulation.

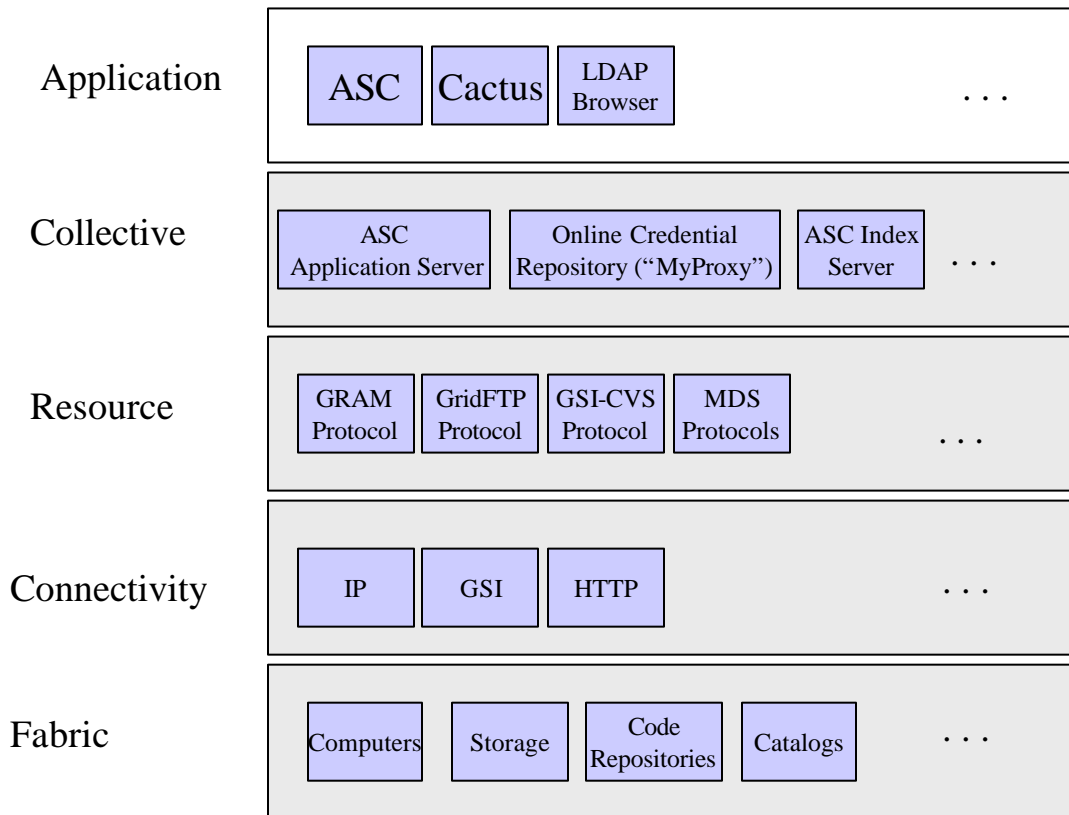


Figure 1: A protocol-oriented view of Grid architecture, showing the principal components used to build the Astrophysics Collaboratory (ASC). Various components, in particular the ASC Application Server, are described in more detail in the text.

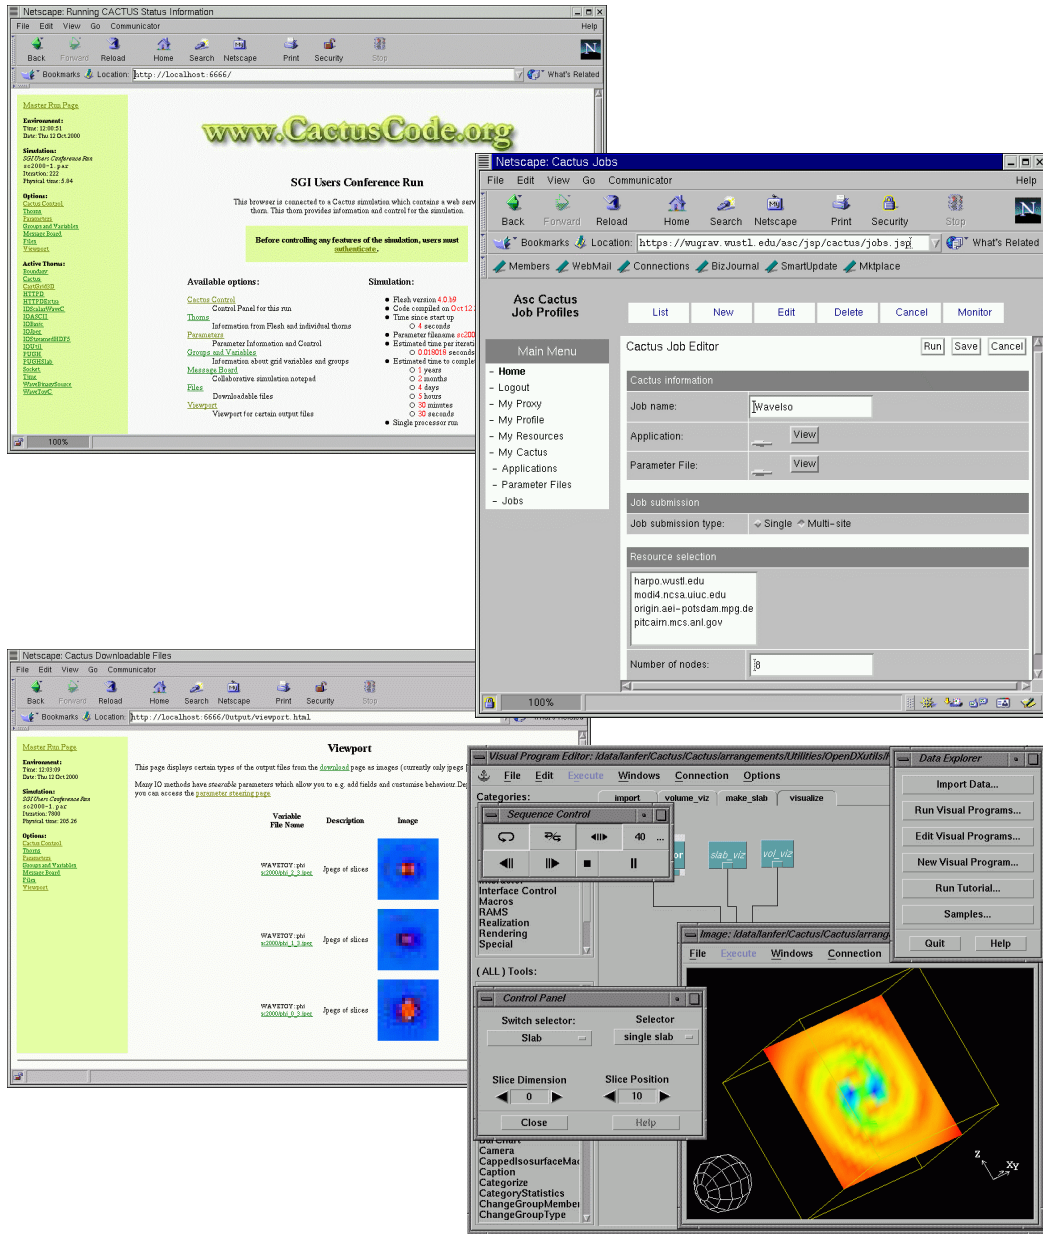


Figure 2: The look and feel of the ASC collaboratory exposed through a Web browser and a customized visualization component.

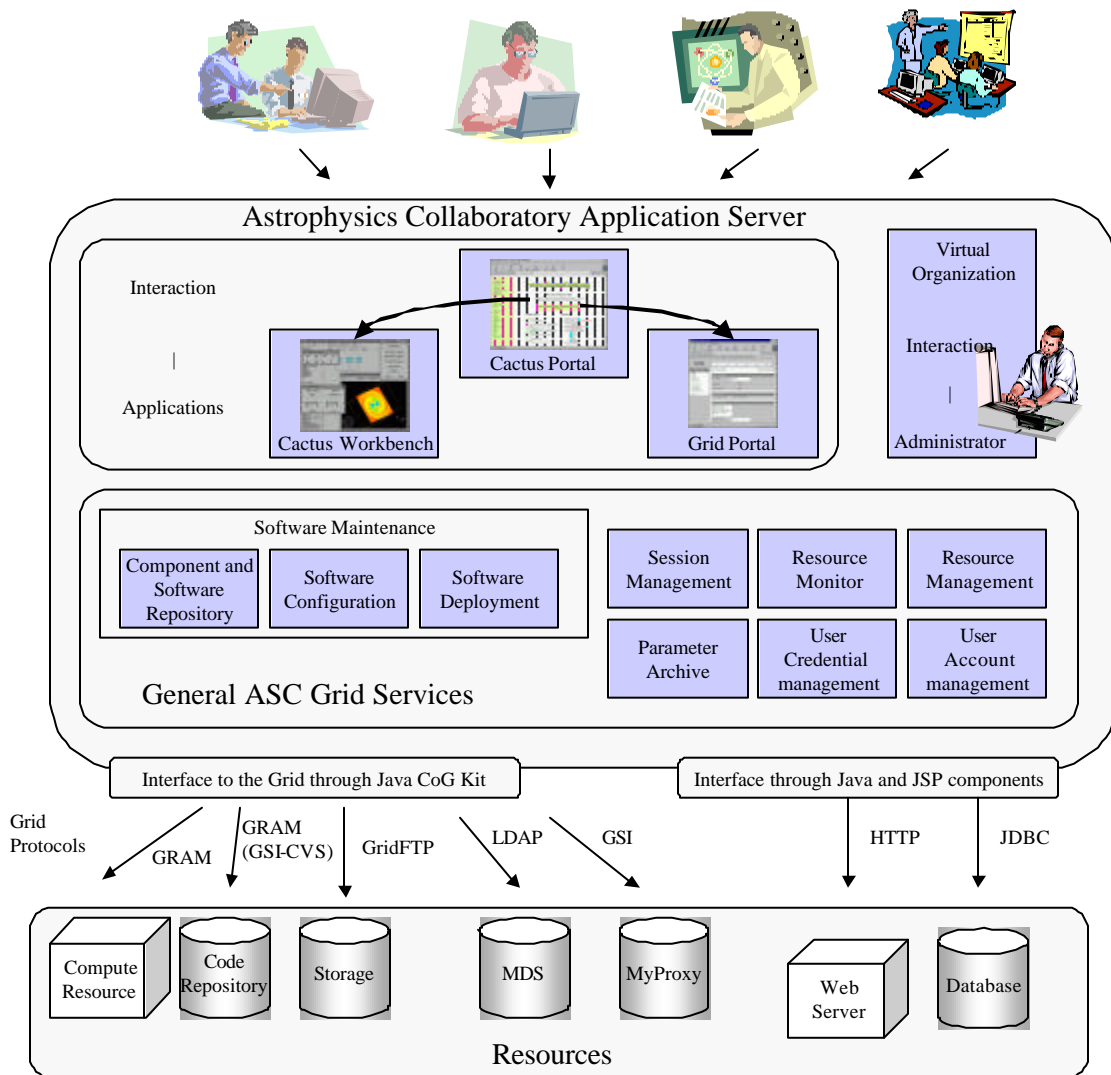


Figure 3: The structure of the ASC Application Server, showing the various components used to meet the requirements identified in the text.

2 ARCHITECTURE OVERVIEW

As illustrated in Figure 2, the ASC has a modular structure and makes heavy use of Grid protocols and services to access and manage distributed resources. The principal components and protocols used are as follows:

- Web browsers serve as the primary end-user interface, serving dynamic HTML (D-HTML) and Java applets obtained from the ASC Application Server. Other end-user applications are also used: for example, for more advanced visualization functions (Figure 2).
- The ASC Application Server is the middle-tier entity that implements most ASC-specific logic, interacting with:
 - Grid services for purposes of resource discovery and authentication;
 - Resources to initiate and manage simulation and/or access and analyze data, and
 - Web browsers and other applications for user interaction.
- An online credential repository, MyProxy [18], is used as a secure cache for user proxy credentials, hence reducing barriers to the use of PKI authentication.
- An ASC Index Server supports registration and subsequent discovery of resources contributed to the ASC, using MDS-2 protocols [9].
- Standard Grid protocols [14] provided by the Globus Toolkit [12] are used to negotiate secure access to remote resources, including CVS code repositories.

The vast majority of ASC-specific logic is incorporated within the ASC Application Server. This component is responsible for handling all interactions with ASC users, for maintaining state associated with user interactions, and for coordinating the use of remote resources to meet user requests. Figure 4 shows its internal structure. Major components include the following

- We make extensive use of commodity Web technologies such as [Java Servlets](#) and JSP.
- We use software development kits (SDKs) and APIs provided by the Java CoG Kit to execute the protocols used to access remote resources.
- We use new SDKs developed for this application (but of general utility) to address issues of code management, deployment, and simulation execution on Grid resources.

The final major component of our collaboratory is the Cactus code[2], a modular simulation development framework that allows us to encourage and leverage the development of reusable components within the astrophysical simulation community.

3 TECHNOLOGIES USED

We use three primary technologies to construct our collaboratory: the Cactus framework, to support development of modular simulation codes; the Globus Toolkit, for Grid services; and various Java and Web technologies, to facilitate construction of the ASC Application Server.

3.1 The Cactus Framework: Modular Simulation and Code Creation

Cactus is a modular simulation code framework [3] that uses component technologies to facilitate the construction and management of complex codes. The Cactus project has established its own terminology for naming different components. They term the code that coordinates the activities of modules the *flesh*, and a module that performs activities and is coordinated by the flesh a *thorn*.

Thorns interoperate via standard interfaces and communicate directly with the flesh, but not with one another. Hence, the development of thorns can be performed independently of one another, which is essential for coordinating component development in large geographically dispersed research groups while being able to integrate specialized thorns developed by a large variety of groups. This separation of the thorn development promotes, for example, the independent development of *physics thorns* to model the black hole or hydro evolutions, and *information technology thorns* that provide access to I/O, parallelism programming frameworks, or performance monitoring. Thus the Cactus component model permits astrophysicists with different levels of IT expertise to participate in collaborative code development. The Cactus framework also supports sharing and cloning of thorns as part of multiple code repositories. Cactus is distributed with tools that support access to Cactus modules that are generally applicable to a variety of problem domains. An example provides a thorn that supports remote visualization and steering of running simulations through an HTTP interface to enable seamless integration with Web-based portals. Specialized components for the astrophysics community are readily available. Furthermore, the Cactus framework is used as part of the specialized service to perform astrophysics simulations. This includes distributing and running Cactus simulations on remote resources and is supported on a wide variety of platforms. We use Cactus as part of

ASC to create a simulation code assembled from various thorns, contribute new thorns through a community process, and execute and monitor them on remote resources.

3.2 Globus Grid Services: Enabling Resource Sharing

By using Grid technologies, in our case Globus [11] middleware, we minimize the difficulties introduced by the tremendous variety of resource types, mechanisms, and access policies, and the complex coordination issues (e.g., relating to authentication and authorization), that we would otherwise have to handle ourselves as part of developing a collaboratory.

The Globus Toolkit [11, 14] is a set of services and software libraries that support Grids and Grid applications. The Toolkit includes software for security, information infrastructure, resource management, data management, communication, fault detection, and portability. It is packaged as a set of components that can be used either independently or together to develop useful Grid applications and programming tools. Globus Toolkit components include the Grid Security Infrastructure (GSI) [6], which provides a single-sign-on, run-anywhere authentication service, with support for delegation of credentials to subcomputations, local control over authorization, and mapping from global to local user identities; the Grid Resource Access and Management (GRAM) protocol and service [10], which provides remote resource allocation and process creation, monitoring, and management services; the Metacomputing Directory Service (MDS) [9, 17], an extensible Grid information service that provides a uniform framework for discovering and accessing system configuration and status information such as compute server configuration, network status, and the locations of replicated datasets; and GridFTP, a high-speed data movement protocol [1]. A variety of higher-level services are implemented in terms of these basic components.

3.3 Java Framework and Server Pages: Developing Browser-accessible Thin Clients

3.3.1 Java Server Pages / Servlets

To enable rapid prototyping of Web portals as part of ASC, we base much of our Web-based services on the JavaServer Pages (JSP) technology [15]. JSP uses XML-like tags and scriptlets implemented in the Java framework to encapsulate the logic that generates the content for the page. We have exposed the functionality of the application server as a set of reusable JSP components that both provide direct access to Grid

services and also implement higher-level services such as monitoring of the compute resources, execution of complex jobs with file staging and cleanup abilities, and maintaining session state between ASC accesses. These components can be reused in other projects.

3.3.2 *JDBC*

Within the ASC Application Server we use a relational database to maintain the state of sessions. We express all database operations in terms of Java Database Connectivity (JDBC) [16] calls, which provide a generic interface to Standard Query Language (SQL)-compliant database management systems. Thus the ASC can support all commercially viable SQL-based servers, all of which are distributed with a JDBC provider. In our current implementation, we use the popular open-source database server MySQL which is free for non-commercial use, well documented, reliable, and provides high performance.

3.3.3 *Java CoG Kit*

Within the ASC Application Server, we use SDKs provided by the Java Commodity Grid Toolkit [22] to access Globus Toolkit services, as follows:

- **Security:** Access to GSI services and to the MyProxy [18] online credential repository for retrieval of certificates.
- **Information Services:** Access to MDS information services through the Java Naming and Directory Interface (JNDI) [20].
- **Job Submission:** Remote job submission and management through GRAM client APIs and protocols. This allows access to Grid resources including batch-scheduling systems on high performance computers.
- **File Transfer:** Secure access to GSI-enabled FTP servers for file transfer [1] as well as access to GASS servers [4].

3.3.4 *Other Collaboratory Technologies: Collaboration through Commodity Technologies*

We intend to integrate other collaborative technologies such as audio/video conferencing, chat rooms, news, and e-mail groups. We have refrained from doing so in this phase of the project because the access to the Grid and the code repository was of most benefit to the user community we targeted first.

4 PORTAL-SPECIFIC SERVICES AND COMPONENTS

In this section we discuss services and components that we have developed to support the Web-based collaboratories. In all cases we rely upon GSI security wherever possible.

4.1 Secure Login

In order to gain access to services provided by the ASC, users logon to the ASC Portal with a secure Web form. Once logged on users may retrieve their GSI proxy certificates from the ASC MyProxy server at any time. From there, users can make use of our application services without the need to supply additional passwords until either the user logs off or their GSI proxy certificates expire.

4.2 User Management

We have developed classes for maintaining information about users of the Collaboratory and the tasks they perform. Users can track the remote jobs they have started for instance, or the file transfers they have performed. Users can also maintain a list of GSI certificates they use to gain access to remote resources and easily determine to which services their certificates successfully authenticate.

In addition, users can be granted or denied access to the functions of our Web-based Collaboratory, in particular for enabling select users to administer the ASC Application Server, including adding or removing users from the Collaboratory, monitoring user sessions and starting or stopping application server tasks.

4.3 Grid Resource Management

We attempt to model computational resources and the services they support in an intuitive way, by providing classes that abstract the details of the MDS queries required to gather information from remote resources and the services available on those resources. Additional information about each resource, such as the location of software not otherwise made available to MDS, can be maintained in the Application Server's primary database and published to users as needed.

4.4 Grid Service Management

As mentioned above, we can track what services are available on remote computational resources, where for each service we attempt to provide intuitive interfaces for invoking, testing, and monitoring user requests to those services. In most cases, we extend upon the JavaCoG Toolkit to allow user activities to be written to the ASC Application Server database and provide wrappers that unify the interfaces to those services.

4.5 Software Management

We support commonly used tools for managing the development of software within the Collaboratory. In particular, we provide interfaces for invoking CVS and MAKE on remote computers with interactive GRAM jobs, thus taking advantage of GSI security. In the case of CVS, we worked with the Globus Project to develop GridCVS. GridCVS extends the CVS protocol and service to include GSI-authentication and we are developing a Java-based client that allows checkouts/checkins to be performed from/to GSI-enabled CVS repositories from third-party locations, where remote file browsing is performed with our GridFTP Java classes, extended from the JavaCoG Toolkit.

4.6 Cactus Management

Since the ASC supports application development with Cactus, we have developed tools that are useful for manipulating Cactus applications and data on remote resources from our Web-based Collaboratory. Some of these tools are described below.

- **Parameter file management** tools to allow users to maintain a set of parameter files important for the execution of an application. A database is used to maintain the various parameter files and share them within the collaboratory.
- **Simulation management** tools to allow executing domain and non-domain Cactus jobs. Cactus jobs can be invoked as GRAM requests or be run from batch scripts that include pre- and post-processing instructions.

4.7 Visualization

We are developing tools that allow users to run applications for visualizing remote datasets. Currently, for those users running X-servers on their local desktop, we provide a Web interface that enables users to search for files with our DHTML-based GridFTP file browser on remote machines and subsequently launch visualization applications that exist on those machines, such as LCAVision, an application developed by the ASC

for visualizing HDF5 data, with the application's display set to the user's local desktop. We have found this works well in practice. However, we are planning on developing additional tools for manipulating remote datasets so that desired portions of remote datasets can be viewed locally as well.

5 IMPLEMENTATION

Our portal can be accessed from any DHTML enabled browsers. We have tested our software with such popular browsers such as Internet Explorer v4 and Netscape v4. The Cactus software runs on almost all major operating systems. A prototype for Windows NT is available. New remote resources are brought into the collaboratory by implementing the Grid protocols and services used for remote discovery, job submission, and data transfer. In general, these ports are provided by the Globus Toolkit. The Web server is based on a secure version of Apache with JSP extensions. Persistent portal session management is achieved via MySQL.




Figure 4: The ASC project investigates astrophysical phenomena involving strong dynamic gravitational fields such as this grazing collision of two neutron stars. The surfaces represent the density of the stellar matter, while the warping of the gridded surface represents the gravitational flux on the plane of the collision.

6 APPLICATION

ASC was originally developed in support of numerical simulations in astrophysics, and has evolved into a general-purpose code for partial differential equations in three-dimensions. Perhaps the most computationally demanding application that has been

attacked with Cactus is the numerical solution of Einstein's general relativistic wave equations, in the context, for example, of the study of neutron star mergers and black hole collisions (Figure 4). For this purpose, the ASC community maintains an ASC server and controls its access through login accounts on the server. Remote resources integrated into the ASC server are controlled by the administrative policies of the cite contributing the resources. In general, this means that a user must have an account on the machine on which the service is to be performed.

The modular design of the Cactus framework and its easy exposure through a Web-based portal, as described in this paper, permits a diverse group of researchers to develop add-on software modules that integrate additional physics or numerical solvers into the Cactus framework.

7 RELATED WORK

Much of the related work is documented in this special issue.

8 FUTURE DIRECTIONS

We plan to improve our collaboratory in several ways, with enhancements to the user interface, advanced visualization support, data management tools, collaboration support, and development of new interaction paradigms.

While testing our framework within the astrophysics users community we have observed limits of our initial thin clients based on DHTML. Therefore, we plan to investigate a wider variety of client-side interfaces including fat client and slender client interfaces. These will also include the use of Java3D as an alternative display framework. In cases where speed is essential, we use fat clients that are traditional monolithic applications, such as the OpenDX visualization tool. In these situations, the ASC Application Server merely acts as a broker for creating direct connections between the client and various Grid services. The application server, in this sense, would serve as a sort of Napster for peer-to-peer interactions between clients and distributed resources.

Another area of development is the creation of robust data management services, such as tools to assist users in finding appropriate storage space on remote computers, archival storage management services, metadata browsing, and automated replication of

data to high-performance disk caches to improve visualization system performance. This data management capability is critical for providing access to remote visualization services as well as sharing and understanding the large volumes of archived results produced by this community of researchers. In the long term we intend to explore alternative modes of interaction with our Grid Portal through asynchronous collaborative methodologies. We also plan to reuse the ASC framework in other application domains, for example in, in high energy physics.

9 CONCLUSION

We have described an online collaboratory framework that supports the collaborative development and use of numerical simulation codes. This framework uses a combination of familiar Web and Java technologies; Grid services for secure access to and coordinated use of distributed resources; and custom components for collaborative code development and distributed simulation. We have described an application of this framework within the astrophysics community, and suggest that the framework can also be applied within other application domains with similar requirements.

ACKNOWLEDGMENTS

The development of an *Astrophysics Simulation Collaboratory* joins 5 institutions (Washington University, Rutgers University, National Center of Supercomputing Applications, University of Chicago, University of Illinois) in an NSF funded *Knowledge and Distributed Intelligence* program. We thank the Numerical Relativity Group at the Max Planck Institute in Potsdam, Germany for developing the Cactus code. This work was supported in part by the National Science Foundation's Knowledge and Distributed Intelligence project "Astrophysical Simulation Collaboratory", and 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. W. Allcock, I. Foster, S. Tueke, A. Chervenak and C. Kesselman. Protocols and Services for Distributed Data-Intensive Science, 2001, <http://www.globus.org/research/papers/ACAT3.pdf>.
2. G. Allen, W. Benger, T. Dramlitsch, T. Goodale, H. C. Hege, G. Lanfermann, A. Merzky, T. Radke, E. Seidel and J. Shalf Cactus Tools for Grid Applications. *accepted for publication in Intl. Journal of Cluster Computing*, <http://www.cactuscode.org/Showcase/Publications.html>.
3. G. Allen, W. Benger, T. Goodale, H.C. Hege, G. Lanfermann, J. Masso, A. Merzky, T. Radke, E. Seidel and J. Shalf Solving Einstein's Equations on Supercomputers. *IEEE Computer*. 52-59, <http://www.cactuscode.org>.
4. J. Bester, I. Foster, C. Kesselman, J. Tedesco and S. Tuecke. GASS: A Data Movement and Access Service for Wide Area Computing Systems. in *Proc. IOPADS'99*, 1999.
5. K. Bishop and J. Alameda, A Prototypical Example of the Chemical Engineer's Workbench. in *AAAS/SURA/NCSA Meeting on Networking Resources for Collaborative Research in the Southeast*, (1998), <http://www.aaas.org/spp/dspp/rcp/gamtg/agenda.htm>.
6. R. Butler, D. Engert, I. Foster, C. Kesselman, S. Tuecke, J. Volmer and V. Welch Design and Deployment of a National-Scale Authentication Infrastructure. *IEEE Computer*, 33 (12). 60-66.
7. R. Cartwright, J. Dennis, R. Hood, J. R. Jump and K. Kennedy Parasol: A Laboratory for Parallel Software Technology Research. (TR86-41).
8. R. Chandra, A. Gupta and J. Hennessy. Data Locality and Load Balancing in COOL. in *Proceedings of the Fourth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, San Diego, CA, 1993.
9. K. Czajkowski, S. Fitzgerald, I. Foster and C. Kesselman, Grid Information Services for Distributed Resource Sharing. in *Proc. 10th IEEE International Symposium on High Performance Distributed Computing*, (2001), <http://www.globus.org>.
10. K. Czajkowski, I. Foster, N. Karonis, C. Kesselman, S. Martin, W. Smith and S. Tuecke, A Resource Management Architecture for Metacomputing Systems. in *Proceedings of IPPS/SPDP '98 Workshop on Job Scheduling Strategies for Parallel Processing*, (1998).
11. I. Foster and C. Kesselman. The Globus Project: A Status Report. in *Proc. Heterogeneous Computing Workshop*, IEEE Press, 1998, 4-18.

-
12. I. Foster and C. Kesselman Globus: A Metacomputing Infrastructure Toolkit. *International Journal of Supercomputer Applications*, 11 (2). 115-128.
 13. I. Foster and C. Kesselman (eds.). *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, 1999,
 14. I. Foster, C. Kesselman and S. Tuecke The Anatomy of the Grid: Enabling Scalable Virtual Organizations. *International Journal of High Performance Computing Applications*, 15 (3). 200-222, www.globus.org/research/papers/anatomy.pdf.
 15. M. Hall *Core Servlets and JavaServer Pages (JSP)*. Prentice Hall PTR/Sun Microsystems Press, 2000,
 16. G. Hamilton, R. G. G. Cattell and M. Fisher *JDBC database access with Java : a tutorial and annotated reference*. Addison-Wesley, Reading, Mass., 1997,
 17. G. v. Laszewski, S. Fitzgerald, I. Foster, C. Kesselman, W. Smith and S. Tuecke. A Directory Service for Configuring High-Performance Distributed Computations. in *Proc. 6th IEEE Symp. on High-Performance Distributed Computing*, 1997, 365-375.
 18. J. Novotny, S. Tuecke and V. Welch, An Online Credential Repository for the Grid: MyProxy. in *to be published in HPDC*, (2001).
 19. *NRC National Collaboratories: Applying Information Technology for Scientific Research: National Research Council*, 1993,
 20. S. S. Rosanna Lee *JNDI API Tutorial and Reference: Building Directory-Enabled Java Applications*. Addison-Wesley, 2000,
 21. M. Russell, G. Allen, G. Daues, I. Foster, T. Goodale, E. Seidel, J. Novotny, J. Shalf, W.-M. Suen and G. V. Laszewski, The Astrophysics Simulation Collaboratory: A Science Portal Enabling Community Software Development. in *Tenth IEEE International Symposium on High Performance Distributed Computing (HPDC10)*, (San Francisco, 2001), <http://www.cactuscode.org/Showcase/Publications.html>.
 22. G. von Laszewski, I. Foster, J. Gawor and P. Lane A Java Commodity Grid Kit. *Concurrency and Computation: Practice and Experience*, 13 (8-9). 643-662, <http://www.globus.org/cog/documentation/papers/cog-cpe-final.pdf>.