ABRIDGED GACP 211–004

A NATIONAL COLLABORATORY TO ADVANCE THE SCIENCE OF HIGH TEMPERATURE PLASMA PHYSICS FOR MAGNETIC FUSION

In response to Office of Science Financial Assistance Program
Notice 01–06 for the U.S. Department of Energy Office of Science

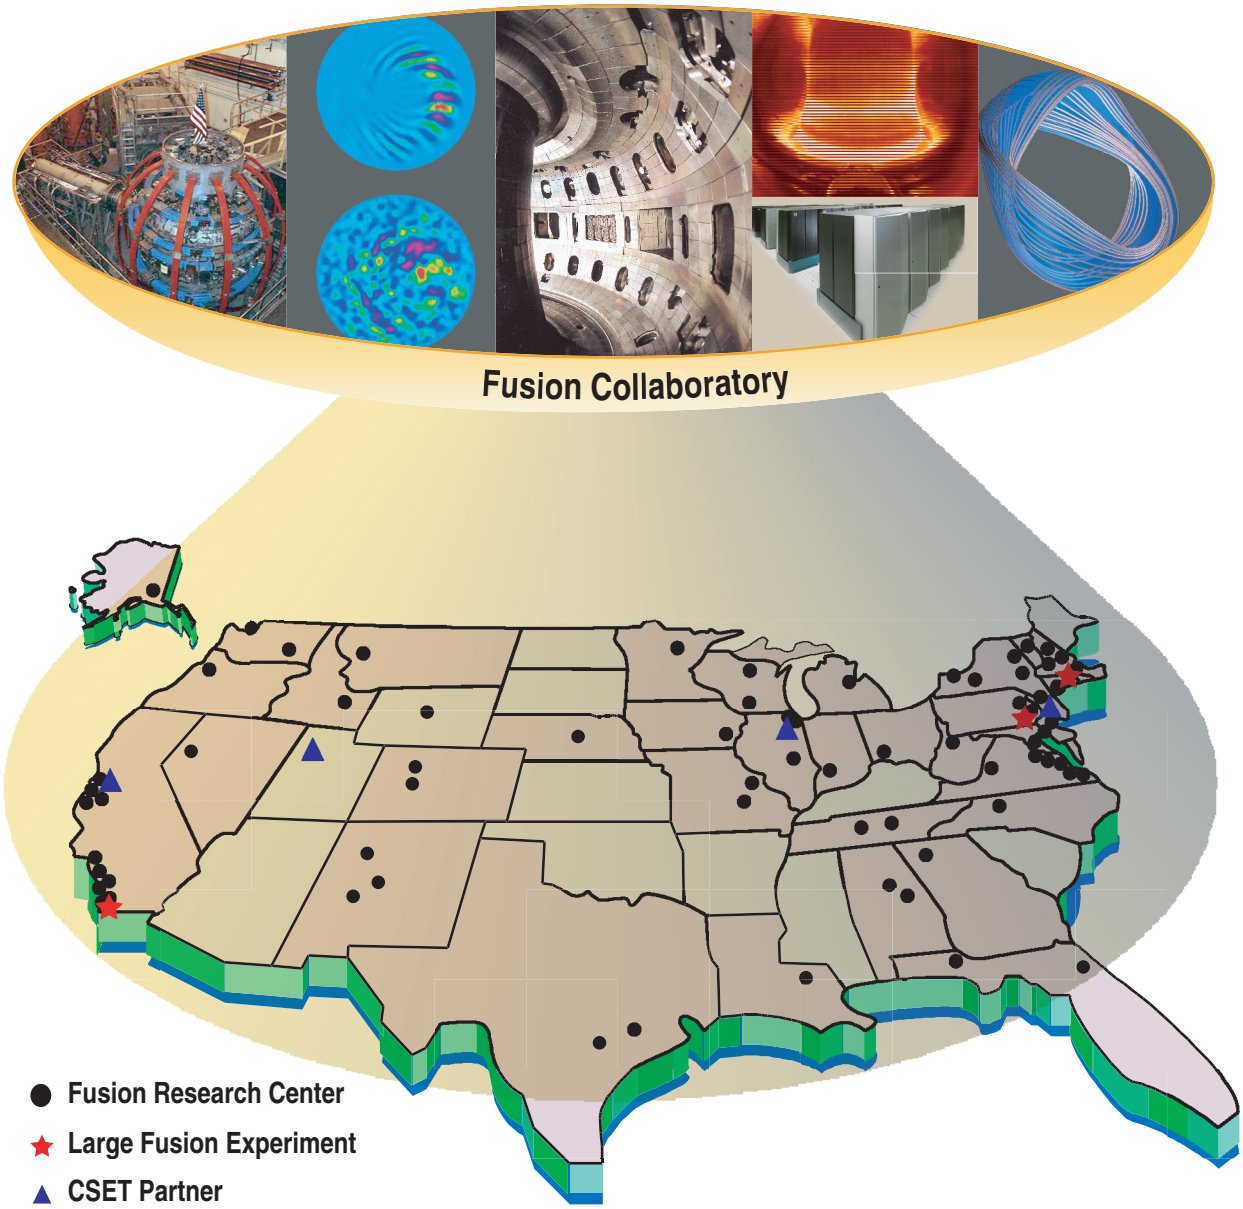

March 15, 2001

ARGONNE NATIONAL LABORATORY · UNIVERSITY OF CHICAGO

MIT PSFC

SCI INSTITUTE

BERKELEY LAB

PPPL PRINCETON PLASMA PHYSICS LABORATORY

GENERAL ATOMICS

GACP 211–004

A NATIONAL COLLABORATORY TO ADVANCE THE SCIENCE OF HIGH TEMPERATURE PLASMA PHYSICS FOR MAGNETIC FUSION

In response to
Office of Science Financial Assistance Program
Notice 01–06 for the U.S. Department of Energy
Office of Science

USE AND DISCLOSURE OF DATA

This proposal or quotation includes data that shall not be disclosed outside the Government and shall not be duplicated, used, or disclosed — in whole or in part — for any purpose other than to evaluate this proposal or quotation. If, however, a contract is awarded to this offeror or quoter as a result of, or in connection with, the submission of this data, the Government shall have the right to duplicate, use, or disclose the data to the extent provided in the resulting contract. This restriction does not limit the Government's right to use information contained in this data if it is obtained from another source without restriction. The data subject to this restriction are contained in all sheets herein.

March 15, 2001

Technical Proposal

EXECUTIVE SUMMARY

The National Fusion Collaboratory will advance scientific understanding and innovation in magnetic fusion research by enabling more efficient use of existing experimental facilities and more effective integration of experiment, theory, and modeling. Specifically, this project will create and deploy collaborative software tools throughout the national magnetic fusion research community comprised of over one thousand researchers from over forty institutions. Built on a foundation of established computer science toolkits, successful deployment of the Collaboratory will nevertheless require significant computer science research to extend the toolkits beyond their present capabilities. The National Fusion Collaboratory will enable networked real–time data analysis and instantaneous communication amongst geographically dispersed teams of experimentalists and theoreticians. This represents a fundamental paradigm shift for the fusion community where data, analysis and simulation codes, and visualization tools will be thought of as network services. In this new paradigm, access to resources (data, codes, visualization tools) is separated from their implementation, freeing the researcher from needing to know about software implementation details and allowing a sharper focus on the physics.

Magnetic Fusion Research. Fusion, the power source of the stars, has been the subject of international research since the late 1950's. Experimental magnetic fusion energy research in the United States is centered at three large facilities with a present day replacement value of over $1B. Teaming with this experimental community is a theoretical and simulation community that concentrates on the creation of realistic 3D plasma models. As the capabilities of wide area networks have increased, more researchers have begun to collaborate with the experimental institutions from their home laboratory rather than travelling for experiments. While the community has made significant progress in accommodating this new pattern of use, much remains to be done to take full advantage of newly emerging technologies.

Goals of the Collaboratory. The aim of the Collaboratory is to:

- Create transparent and secure access to local/remote computation, visualization, and data servers.
- Develop collaborative visualization that allows interactive sharing of graphical images among control room display devices, meeting room displays, and with offices over a wide area network.
- Enable real–time access to high–powered remote computational services allowing such capabilities as between pulse analysis of experimental data and advanced scientific simulations.

Magnetic fusion experiments operate in a pulsed mode producing plasmas of up to 10 seconds duration every 10 to 20 minutes, with 25–35 pulses per day. For each pulse up to 10,000 separate measurements versus time are acquired representing hundreds of Megabytes of data. Decisions for changes to the next plasma pulse are made by data analysis conducted within the roughly 15 minute inter–pulse interval. This mode of operation places a large premium on rapid data analysis that can be assimilated in near–real–time by a geographically dispersed research team.

Benefits to Fusion. The National Fusion Collaboratory will increase physics productivity by

- Enabling more efficient utilization of experimental time on the three large facilities through more powerful between pulse data analysis resulting in a greater number of experiments at less cost.
- Allowing more transparent access to analysis and simulation codes, data, and visualization tools resulting in more researchers having access to more resources.
- Creating a standard tool set for remote data access, security, and visualization allowing more researchers to build these services into their own tools.
- Facilitating the comparison of theory and experiment.
- Facilitating multi–institution collaborations.

The Collaboratory will also increase the productivity of code and tool developers by

- Supporting SciDAC computational initiatives under the collaboratory framework.
- Supporting more users with fewer installations at reduced cost.
- Facilitating shared code development projects resulting in more rapid code creation.
- Creating a standard tool set for remote data access, security, and visualization allowing these services to be easily built into new tools.

Computer Science Research. To accomplish these goals, fusion scientists with expertise in large experiments and simulation code development will join computer scientists with expertise in security, distributed computing, and visualization to form a closely coordinated team. Within each group, there is a long history of collaborative research and our management plan will insure the unified operation of the multi–disciplined groups. This team, leveraging existing computer science technology where possible, will deploy a collaboratory prototype. For requirements not met by current capabilities, new technologies will be developed. The variety of users, resources, applications, and policies encountered will serve as an excellent proving ground for new technologies that prepares the way for their use in other scientific disciplines.

Benefits to Computer Science Toolkits. The Collaboratory will enhance existing toolkits by

- Enabling automatic propagation of security credentials in multi–server contexts.
- Allowing complex use policy implementation and remote monitoring of computational resources.
- Extending security architecture to encompass commercial databases.
- Creating pre–emptive scheduling capability and advance reservation of computational resources.
- Extending the Access Grid to large tiled display walls.
- Developing a quantitative visualization capability allowing data comparison with uncertainties.
- Supporting numerous SciDAC computer science projects that cross–cut with the Collaboratory.

The computer science research necessary to create the Collaboratory is centered on three main activities: security, remote and distributed computing, and scientific visualization.

Security. The sharing of data, codes, and visualization tools as network services requires a system for protecting these valuable resources against unauthorized use. The Collaboratory will exploit state-of–the–art authentication, authorization, and encryption technologies provided by the Globus Security Infrastructure and the Akenti authorization service. Existing fusion community codes will be modified to use this infrastructure for remote execution and data access. To meet the needs of the Collaboratory, the current version of these middleware tools will be extended. In particular, it will be necessary to incorporate rules for fair use of shared resources into the policy enforced by the security model and to enhance tools that enable valid credentials to propagate automatically from resource to resource.

Distributed Computing. The remote and distributed computing requirements of the Collaboratory will utilize the Globus facilities including remote job scheduling, monitoring, exception handling, and accounting. This will enable researchers and their institutions to share the community's computational resources. The components of the Globus toolkit that can be immediately deployed to create the foundation of the Collaboratory are Grid Information Services, Grid Security Infrastructure, and Globus Resource Allocation Manager. Research components that are required to fully meet the needs of the Collaboratory and that will create new functionality for Globus include managing batch versus preemptive job priorities, providing status display and accountability to users, monitoring the adherence of resources to policies, and providing advance reservations. Fusion community codes to be adapted to the Collaboratory include serial and parallel MHD stability codes and serial and parallel transport codes.

Scientific Visualization. The demand placed on visualization tools by the Collaboratory is intense due to both the highly collaborative nature of fusion research and the dramatic increase in data resulting from the enhanced computing capabilities. The visualization component of the Collaboratory will focus on the development of a collaborative control room, collaborative meeting room, and enhanced visualization tools. The collaborative visualization requirements will utilize the Access Grid that enables distributed meetings and collaborative teamwork sessions. Extensions to the Access Grid software include a more closely integrated shared experience with the researcher's current work environment and the support of large tiled displays that will provide collaborative capabilities to large-format remote visualizations. New software will be tested on display walls that already exist within the proposal team. Extensions to display wall software include the ability to have visualizations not tied to an individual projector allowing the size of the visualization to vary depending on the researcher's need. Extensions to the visualizaton toolkit will be the ability to quantitatively compare theory to experiment with uncertainties.

Project Scope. The project will be a three–year collaboration between experts in the magnetic fusion and the computer science and enabling technology communities. The proposed collaborative effort will cost $2M per year with two–thirds of the funds allocated to the computer science research partners.

1. ABSTRACT

The long–term vision of the National Fusion Collaboratory described in this proposal is to transform fusion research and accelerate scientific understanding and innovation so as to revolutionize the design of a fusion energy source. The project will create and deploy collaborative software tools that will enable more efficient utilization of existing experimental facilities and more effective integration of experiment, theory, and modeling. These tools will link together the large magnetic fusion research community that is spread over more than 40 sites in 37 states, enabling networked real–time data analysis and instantaneous communication amongst geographically dispersed teams of experimentalists and theoreticians. Such a collaboratory represents a dramatic shift for the fusion community where access to resources (data, codes, visualization tools) would be separated from details of their implementation and eliminate barriers to their widespread use by the research community (Fig. 1).

The team submitting this grant application is comprised of scientists and software engineers from the three large experimental magnetic fusion facilities and a group of computer scientists from universities and national laboratories with the expertise to achieve the research and development advances required to make the Collaboratory a success. The computer science research necessary to create the Collaboratory is centered on three main activities: security, remote and distributed computing, and scientific visualization. This research will enhance existing computer science toolkits and have a cross–cutting benefit to other components of the SciDAC initiative. The large scale and unique characteristics of the Fusion Collaboratory make it an important testbed for these toolkits. The proposed collaborative project spans a three–year effort costing $2M per year with two–thirds of the funds allocated to computer science research.

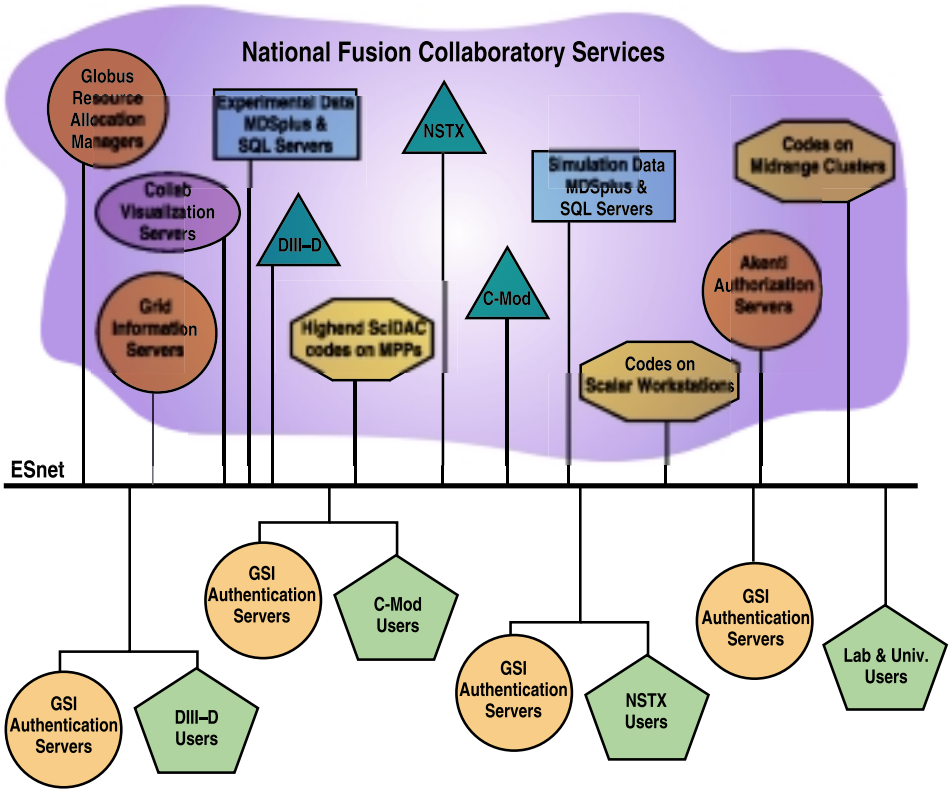

Fig. 1. The National Fusion Collaboratory will provide a unified framework in which data, codes, and visualization tools are available securely and transparently over the Internet.

Contact: D.P. Schissel, fax 858-455-4156, phone 858-455-3387, e-mail schissel@fusion.gat.com

2. NARRATIVE

2.1. Background and Significance

2.1.1. Background

The long–term goal of fusion research has been the development of a reliable energy system that is environmentally and economically sustainable. Nuclear fusion, the power source of the stars, has been the subject of international research for over four decades. To achieve this goal, it has been necessary to develop the science of plasma physics, a field with close links to fluid mechanics, electromagnetism, and nonequilibrium statistical mechanics. Recognizing the potential of fusion energy, the President's Committee of Advisors on Science and Technology (PCAST) designated fusion as an integral component of the nation's long–term research portfolio for national energy security and climate change remediation [1]. Experimental magnetic fusion research in the US is centered at three large facilities (Alcator C–Mod [2], DIII–D [3], NSTX [4]) with a present day replacement value of over $1B. Teaming with this experimental community is a theoretical and simulation community that concentrates on the creation of realistic non–linear 3D plasma models. Working together to advance scientific understanding and innovation, these two groups represent over one thousand scientists from over forty institutions. As the capabilities of wide area networks have increased over the past decade, more researchers have begun to collaborate with the experimental institutions from their home laboratory rather than travelling for experiments or permanently relocating. While the community has made significant progress in accommodating this new pattern of use, much remains to be done if we wish to take full advantage of newly emerging technologies.

Magnetic fusion experiments operate in a pulsed mode producing plasmas of up to 10 seconds duration every 10 to 20 minutes, with 25–35 pulses per day. For each pulse up to 10,000 separate measurements versus time are acquired at sample rates from kHz to MHz, representing hundreds of Megabytes of data. Decisions for changes to the next plasma pulse are informed by data analysis conducted within the roughly 15 minute between-pulse interval. This mode of operation places a large premium on rapid data analysis that can be assimilated in near–real–time by a geographically dispersed research team.

Although the fundamental laws that determine the behavior of fusion plasmas are well known, obtaining their solution under realistic conditions is a scientific problem of enormous complexity. This is due in large part to the enormous range of temporal and spatial scales involved and is therefore a challenge that can only be met with advanced scientific computing. Because of this complexity, researchers have a long history of productive use of advanced computation and modeling. Such work is typically broken down into large (macro) and small–scale (micro) disturbances that relax the plasma to a lower energy state. Both macro and micro scale computing problems have been funded by the Department of Energy (DOE) OFES Plasma Science and Advanced Computing Initiative (PSACI) and are being proposed under SciDAC.

The magnetic fusion community's desire for more efficient collaboration is well known and has been most recently identified in a review by the National Research Council [5]. Efforts to improve collaboration within the community have included sharing of resources and co–development of tools but these have been done in mostly an ad hoc manner. The community has considerable experience in placing remote collaboration tools into the hands of real users. The ability to remotely view operations and to control selected instrumentation and analysis tasks was demonstrated as early as 1992 [6]. Full remote operation of an entire tokamak experiment was tested in 1996 [7,8]. Today's experiments invariably involve a mix of local and remote researchers. Support for remote collaborations was facilitated by the adoption of a common data acquisition and management system, MDSplus [9] and common relational database software Microsoft SQL Server [10]. MDSplus was developed jointly by Massachusetts Institute of Technology (MIT), Los Alamos National Laboratory (LANL), and the Istituto Gas Ionizzati in Padua, Italy. Based on a client/server model, MDSplus provides a heirarchical, self-descriptive structure for simple and complex data types. It is currently in use on 14 experiments around the world storing digitized, analyzed and simulation code data.

2.1.2. Concept

The National Fusion Collaboratory represents a fundamental paradigm shift for the fusion community where all data, analysis and simulation codes, and visualization tools will be thought of as network services. In this new paradigm, the use of resources (data, codes, visualization tools) is separated from their implementation, freeing the researcher from the need to know, in detail, how resources are implemented.

2.1.3. Goals

The proposed National Fusion Collaboratory aims to catalyze a dramatic change in the way research is carried out by enabling the accelerated development of computational tools and techniques so as to revolutionize the understanding and design of a fusion energy source. We propose to eliminate many of the barriers that currently restrict effective collaboration among the geographically diverse fusion research community. The three large experimental facilities will be more efficiently utilized through enhanced collaborative data analysis during machine operation. At the same time, a more efficient diverse development and usage environment will be created for large–scale analysis and simulation codes. Taken together, the Collaboratory will allow for more effective integration of experiment, theory, and modeling on a broader scale. In the process, we will also advance the state of the art in collaboratory technologies via a process of integration, development and evaluation. While providing strong support for the fusion program, very little of the work is discipline-specific. The computer science development and the results of deployment work will be directly applicable to other fields, such as chemistry, climate, and materials science, that require community access to codes and data.

2.1.4. Anticipated New Research Environment

The following usage scenarios illustrate types of collaboration that are presently not possible but that have been requested by the fusion community.

- Enhanced Experimental Operations: In the first minute after a DIII–D plasma pulse, a scientist based at Wisconsin dispatches a detailed physics analysis of the DIII–D data to a Princeton Plasma Physics Laboratory (PPPL) compute farm. This analysis code automatically retrieves data from the DIII–D MDSplus server and saves its results back to the server upon completion. The physicist at Wisconsin generates a visualization comparing the results to other analysis run locally at DIII–D. The visualization is shared and discussed with the DIII–D control room and with collaborators at Lawrence Livermore National Laboratory (LLNL) and PPPL. Based on this information, the experimental team is able to make a decision that affects the configuration of the following plasma pulse.
- Remote Data Access and Display: An investigator at Columbia logs onto her workstation and authenticates herself. She then runs an analysis code that compares data from General Atomics (GA) and PPPL, displaying the results in a visualization that can be seen simultaneously by researchers at all three sites. Spotting a particular regime of interest, she has the data processed by a powerful rendering engine at GA and an interactive 3D display is made available to all meeting participants.
- Remote/Distributed Analysis: A researcher at MIT acquires data from a local experiment, preprocesses the data on local computers, creates an entry for a code run in a local relational database then sends a message to a queue server at PPPL to run a large analysis job. The queue server verifies his identity, finds that he is authorized to use the code and compute facilities and queues the job on a compute cluster. The job reads data over the Internet from MIT, passing along the credential that identifies him as an authorized user, does its computation and writes the results back into the MIT data archive. It also reaches over the Internet to update the relational database entry to reflect the successful completion of the job.
- Detailed Comparisons Between Simulation Data and Experiments: A principal investigator (PI) from LLNL runs a very large simulation on an MPP at NERSC [11], writing the results onto a data server at GA. He runs a synthetic diagnostic on the data on his local workstation that accesses the necessary portions of the simulation run over the net. He then pulls up fluctuation data from the experiment at MIT that has been taken with a physical diagnostic with the same characteristics as the one run on the simulation. The two data sets along with the experimental uncertainties are compared and displayed.

2.1.5. Technical Approach

We will address the technical requirements just listed via the creation, deployment, operation, and evaluation of the National Fusion Collaboratory. This Collaboratory will join the more than one thousand members of the fusion research community by linking resources nationwide, including:

- Computational resources at fusion research sites across the country. The project will start with clusters and workstations at participating institutions, but expect these resources to grow to encompass thousands of processors at the 40+ fusion research sites in the U.S.
- Specialized computational resources such as supercomputers at NERSC. The project will aim to exploit the DOE Science Grid capabilities [12], the subject of a complementary Collaboratory Pilot.
- Data resources across participating institutions, accessed via MDSplus and SQL servers.

In constructing this Collaboratory, we will build extensively on Grid concepts and infrastructure [13,14], exploiting, and in some cases extending, the following components:

- *Security services* enabling “single-sign-on” authentication of Collaboratory users and subsequent authorization during access to Collaboratory resources. The Globus Grid Security Infrastructure (GSI) [15] and Akenti [16] system will be exploited, potentially building on new results and mechanisms proposed [17] by the ANL-ISI-UW SciDAC Collaboratory proposal “Security and Policy for Group Collaboration." GSI will be extended to embrace fusion data systems as well as innovating in the area of authorization.
- *Resource management services* enabling on-demand access to computational resources across the Collaboratory, including NERSC and other DOE resource sites, subject to local and global policies. The project will build on Globus services for resource discovery [18], resource access [19], and computation management [20], and innovate in the area of policy-driven scheduling.
- *Visualization services* enabling collaborative visualization of large and complex datasets produced by fusion simulations across the U.S. Here, we will exploit a considerable experience and technology base at Utah and Argonne National Laboratory (ANL), as well as new technologies to be developed in this project.

2.1.6. Relation to SciDAC and Other DOE Projects

It is anticipated that the proposed project will have a significant impact on the practice of both fusion and computer science research. The PIs are leaders, respectively, in both the fusion, collaboratory/Grid, and visualization communities, and have a strong track record of introducing and gaining acceptance for new technologies. The project will form close connections with a number of existing DOE–funded technology and application projects, as well as with new projects proposed under the SciDAC program. In brief, this project will complement a range of projects that are collectively concerned with developing, deploying, applying, and evaluating a DOE Science Grid Collaboratory Software Environment.

Some of the major connections are:

- The proposed “DOE Science Grid” Collaboratory Pilot will prototype, evaluate, apply, and refine an operational Grid infrastructure across DOE laboratories. The DOE Science Grid project looks to application projects such as the Fusion Collaboratory as their primary source of requirements and evaluation mechanism; in turn, the DOE Science Grid project provides us with an effective deployment mechanism, ensuring for example that GSI enabled servers are deployed at DOE laboratories.
- The submitted "PKI and Directory Services" proposal from ESnet would establish DOE Certificate Authority and information services, hence providing essential infrastructure elements required in this project.
- SciDAC technology proposals [17] from ANL and Lawrence Berkeley National Laboratory (LBNL) will contribute key technology elements concerned with security, policy, and data access.
- We are working with SciDAC fusion sciences application proposals "The Plasma Microturbulence Project," "Center for Extended Magnetohydrodynamic Modeling," "Numerical Computation of

Wave–Plasma Interactions in multi–dimensional Systems" with the aim of incorporating their codes into the collaboratory framework.

- The work described here will build on results obtained at ANL (Globus) and at LBNL (Akenti) and the Diesel Combustion Collaboratory [21] under DOE2000.
- Development of shared and collaborative tools sponsored by OFES will be the basis for much of the user side software incorporated into this proposal. Examples include, MDSplus, the software sharing and remote participation tools developed under the Virtual Laboratory for Technology [22], and the code modules library of the National Transport Code Collaboration [23].
- Computer science technology transfer to international standards bodies [24–26].

2.2. Security

2.2.1. Introduction

Motivation and Overview. In the Collaboratory, data, codes, visualization tools, and experimental instruments can be thought of as network services. At the same time, they represent valuable resources that must be protected against unauthorized use. Complex policies must be embodied in the security model since the resources are typically in different administrative domains with differing rules for access and fair use. To provide security for the Collaboratory, the project will deploy the Globus Grid Security Infrastructure (GSI) [15] and the Akenti authorization service [16] along with the other Globus [27] services such as GRAM [19] and GridFTP [28] that support remote resource access. Existing fusion community codes will be modified to use this infrastructure for remote execution and data access. These middleware tools will be extended to meet Collaboratory needs, for example, to embody fair use of shared resources into the policy enforced by the security model and to enhance tools that enable valid credentials to propagate automatically from resource to resource.

Deployment of Existing Prototype Security Infrastructure. The Collaboratory will exploit state-of-the-art authentication, authorization, and encryption technologies provided by the GSI and Akenti authorization system to construct a highly secure infrastructure. Existing tools in the fusion research community will be fitted with a common user authentication mechanism. While the existing tools have built-in remote access capabilities, they do not have a common mechanism for user authentication and the mechanisms used for resource authorization do not use up-to-date security technology. Deployment of the new technology will greatly improve the security for remote resources and provide an improved and scalable mechanism for managing access controls. At the same time, the production use of these technologies by a large community will provide invaluable feedback to their developers, and will also help guide development of the emerging DOE science grid. The variety of resources, users, and policies encountered in the Collaboratory will demand sophisticated authorization and policy technologies, and hence will serve as an excellent proving ground for new technologies proposed by the ANL and LBNL groups in this area [17].

The current version of GSI allows identity impersonation (or proxying) some fixed pre-set number of levels that facilitate single sign-on and access to secure resources by remote jobs. It allows trust relationships to be statically configured on a per-site basis. Akenti allows remote stakeholders to set static access policies on a per resource level. For example, a user may wish to run a particular code on a compute cluster using data from an experiment's archive as input. Owners of the code, cluster, and data each should be able to set their policies independently and only users with authorization from all three should be able to carry out the task.

Within the fusion community there are many important production codes shared by a geographically diverse users group. There is an ongoing effort to convert all of these codes to use a common data format supported by the MDSplus data servers, and a commercial relational database, Microsoft SQL Server. Globus gateway servers running GRAM and communicating over GSI can provide a secure common interface to the execution code. The MDSplus server and clients are available as source code and will be modified to use GSS/GSI secure communication. Secure, credential based communication with the SQL Server will be accomplished by creating a GSI–enabled gateway server on the SQL Server site that will listen for authentication requests. After using the PKI credential [29] to verify a user, this server will set the user's database password to a random one-time use value. This value is returned through the secure connection to

the client, which can then use it with the standard SQL Server libraries to communication with the SQL Server.

Research Issues. It is anticipated that the fair use of shared compute resources must be provided for during the production stage of this Collaboratory. This new capability will require a policy to define allotments for users, groups, or projects for each of the resources. It will also require accounting procedures to track resource use and an authorization agent that can compare the requested resource use to the allotment less the used amounts to decide if the user should be allowed to schedule a resource. Advanced scheduling and possibly co-scheduling of multiple resources may be required in order to do between-pulse processing during magnetic fusion experiments. This adds a requirement to securely claim a pre-scheduled resource.

Experience in previous distributed computational environments has shown the necessity of allowing programs or active agents to operate on strongly protected resources on behalf of a user. The simple one or two level impersonation credential supported by the current version of Globus are a good start, but considerable research and development needs to be done to design secure, restricted, multi-level delegation that the original delegator, the intermediate servers and the end resource can trust. Some of these issues are proposed in the SciDAC project "Security and Policy for Group Collaboration"; here we propose to extend and apply those techniques to meet the specific requirements of the fusion Collaboratory.

2.2.2. Security Infrastructure

Certificate Authority Deployment. Phase 1 — Use existing certificate authority (CA) to issue certificates. In the first phase the CA running at LBNL as part of the DOE Science Grid can be used to issue certificates to collaboratory users and servers. This will allow deployment of Globus services to the fusion community machines in parallel with the creation of a Certificate Management System (CMS) dedicated to the Collaboratory.

Phase 2 — Deploy a CA specifically for the Fusion Collaboratory. We plan to purchase and run iPlanet [30] Certificate Management System version 4.2. This system provides a Java and Web based interface both to users of the system and administrators. It is designed to allow multiple people at remote sites to issue certificates. It provides for separate registration authorities to give different institutions more control over issuing certificates. Alternatively all certificates can be issued from the CA, but by agents from different organizations. The CA should be run on a dedicated secure machine; one that does not have many user accounts or network services running on it. Ideally, it should be in a location with limited physical access like a locked room.

Server–Side Tools to Enable Policy-Based Resource Management. We plan to add Akenti authorization to GSI. Currently GSI accepts a user's authenticity by checking her identity certificate with a statically configured CA-signing-policy file to see if a user's identity certificate is signed by a CA trusted to sign names of that pattern. If a user passes that test GSI then checks with a static map file to see if this identity maps to a local account. Any actions on behalf of this user are then run under that local account. This model requires that each user have a local account on each machine to which she needs access. The CA-signing policy file is equivalent to Akenti's Root policy file. Both specify the minimum requirements of an identity certificate to grant any access to the resources at a site. Adding Akenti authorization calls to the two methods of data access, MDSplus and the GSI interface to SQL Server, is quite straightforward. In each case the access is being mediated by a server, which can call the Akenti check-access interface to see if the action is allowed.

In the case of executing a job the Akenti check-access interface can return either an individual local user id, a user id usable by a project group, or simply the information that the user has permission to run jobs on this machine. The gateway server who is going to start the job, now has a choice of what user id to run the job under. If the job is going to make file accesses to local files whose access is restricted to specific user ids, it will need to be run as specific id. If it is accessing MDSplus data it may need a delegated credential to allow such accesses.

Refining Akenti User Interface Tools. One of the major problems with the current ad hoc access control mechanisms is that each machine or server has a different method of identifying users and setting access. GSI provides a common way of identifying users. The Akenti User Interface tools provide a com-

mon way for stakeholders, acting remotely, to set access policy for resources. The existing tools consist of client Java GUI applications that communicate with a simple Java server on the resource machine. These interfaces provide the stakeholder with sets of menus from which to choose resources, users, attributes and allowed actions. They then create the digitally signed certificates that are used by Akenti. There are some basic Web-based interfaces that allow a stakeholder or user to view the policy that exists for a resource.

These tools need to be expanded to create the root policy files that are currently just kept as protected text files on the resource server. They will also need to keep pace with the changes required by the new policy demands of the National Fusion Collaboratory. The interfaces to display existing policy also need to be expanded.

Credential Delegation. The current Globus proxy-init mechanism allows a user to type the password for his long-term identity credential once, where it is used to create a short term proxy with an unencrypted private key. This proxy can be used without the user's password for a short-term representation of the user. One of the implications of using a proxy is that the authorizing process on the server machines needs to be able to recognize a proxy and know what user identity it represents. The design proxy certifications are currently being specified in GridForum and IETF drafts [25,26]. Both the current GSI and Akenti will need to be modified to support this format. Most likely the GSI library will provide routines to return the needed information from a chain of CA, User and proxy certificates.

Once a job is executing on behalf of a user, it may need a proxy for which it holds the private key so that it can access the user's protected data. In this case, the agent that started the job may need to get a proxy certificate and put the credential someplace that the job can find it.

2.2.3. Adapting Fusion Applications to use Security Infrastructure

MDSplus Data Access. MDSplus is a data acquisition and data management system used widely in the fusion research community. The remote data access capabilities in MDSplus allow a scientist to access fusion data residing anywhere on the Internet. MDSplus can also be used as a gateway to data stored in legacy data systems. This remote access capability and common data interface has made MDSplus an important tool for sharing data, analysis codes and visualization applications across the entire community. As the use of MDSplus grows it is becoming more evident that more modern techniques will be necessary for providing access control to the data. The current access control functionality provided in MDSplus is quite primitive. MDSplus uses a simple mapping of remote-username and internet address to local accounts on the data server. The connection of the client is trusted without any authentication or password control. The user mapping is managed independently on each server by editing a text file containing the mapping information on each server. Replacing or enhancing the existing access control functionality with GSI and Akenti will greatly improve the security and manageability of resource access control.

The current design of the MDSplus client/server communication should lend itself readily to adding the GSI secure authentication handshaking. When a client wants to connect to an MDSplus server the application issues an mdsconnect call specifying the host and port number to connect to. MDSplus creates a TCP/IP socket and connects to the remote host and sends the username of the client to the remote host. The MDSplus server on the remote host accepts this connection and reads the username from the socket. The server then searches for a mapping between the remote-user@remote-ip-address and a local user in the mapping file. If a valid mapping exists the server switches its identity to the mapped local user. At this point normal user based access control is used on the server system. If the server cannot find a valid remote user to local user mapping, the server responds to the client with an access denied message and the server drops the connection to the client.

This user–based access control which uses file system access control provided by the various operating systems is probably sufficient for remote data access. The parts of the existing MDSplus access control scheme that are inadequate are the remote user authentication and the management of the remote user to local user mappings. Incorporating GSI and Akenti into this mdsconnect dialog will greatly improve these functions. After the client first opens a socket connection to the MDSplus server, this socket can be used by GSI to provide secure authentication to the server. Likewise on the server, when it receives a connection request from a client, the server can pass the connection socket to the GSI to perform a secure dialog with

the client to obtain the identity of the remote client. The server can then use Akenti to determine what resources the remote client should be permitted access.

SQL Server Data Access. Relational databases are crucial for operation of experiments, analysis of data, and for tracking the runs of major physics codes. Unlike MDSplus, the source code for these inexpensive, commercial products is not accessible. Therefore the proposed design provides for secure communication with a commercial database, using the standard communications libraries without modification. This step represents an important extension of the security architecture.

A server on the machine where the database resides (WinNT/2000) will handle secure authentication for SQL Server. This server will authenticate the user using the GSI/Akenti infrastructure and set the user's database password to a random one-time use value. This will be sent back to the client over a secure channel, and the client will then use it to authenticate himself to the database. A timer in the server will be used to reset the password after a short time. The actual communication between the client and the database will be wrapped within the GSI connection.

- The GSI toolkit will be installed locally on a currently supported platform. This will serve as a reference installation for understanding, testing, and debugging GSI applications onsite. The platform will be Linux. Certificates issued by an existing DOE certificate authority, or a locally established CA will be used with this demonstration system.
- The GSI server toolkit will be ported to Windows 2000 so that it can be used with SQL Server database on that platform. After the basic tools and communications infrastructure has been ported the server(s) will be installed as native Windows 2000 services.
- A one-time use password scheme will be implemented for SQL Server. Clients will connect to a GSI server on the system that houses the database. This server will authenticate the user, and determine the mapping to a local user account. It will then reset the user's password to a randomly generated string, and return that string to the client over the secure channel. It will reset the user's password after a predetermined time, on the order of a few seconds, to a new random value. During this time window, the user will authenticate themselves to the database using these temporary credentials.
- All communications between clients and the database will be done using the standard library calls without modification. They will be secured using a secure port–forwarding scheme. The client application will interact with a server that resides locally on the same computer as the client. This server will send the client requests, and receive the server's responses from a server running on the machine with the database, over a secure channel. This server on the database machine, on receipt of requests from the client's secure server, will pass them on to the standard database server application. It will forward all responses from the database back to the client over the secure channel.

The combination of secure communications with one-time use passwords will provide robust, secure access to our existing relational databases. This solution does not require any modification to the database software or network communication library, which is critical since the source code for these is not available. The one-time use passwords will prevent malicious users from gaining access to the database through unsecured channels. The encryption will provide an extra measure of security and prevent IP masquerading attacks.

2.3. Remote and Distributed Computing

2.3.1. Overview

The Fusion Collaboratory will advance fusion science by implementing wide area access to powerful fusion data analysis and simulation codes running as network services. Fusion codes will run on hardware appropriate for each particular code ranging from fast serial workstations to midrange parallel clusters to massively parallel supercomputers. In addition to increasing fusion scientists' capability for offline data analysis, the Fusion Collaboratory will revolutionize fusion experimental physics by enabling detailed time dependent plasma transport and stability analysis *between pulses*.

A major research theme in current experimental plasma physics is the balance between the driving processes and stabilizing processes affecting growth rates of plasma MHD instabilities and turbulent parti-

cle and energy transport, which in turn determines the quality of the fusion plasma. Accurate assessment requires time–dependent plasma transport and confinement analysis, both to validate measurements, and to make a complete and accurate determination of plasma composition and pressure. There follows a detailed analysis of the plasma nonlinear dynamics. The computations involved explain many features of the experimental data but can take several days to execute using the tools currently available to scientists (typically a few workstations running serial codes on a LAN). This turnaround time *can be reduced from days to minutes* by putting the power of a national network of computational services directly in the hands of the fusion community, with a profound benefit for plasma physics and nuclear fusion research.

For the computer scientist, the challenge is to provide an economical unified framework for on demand "over the network" computing to a diverse national collaboration involving three major experimental sites and about a thousand researchers at forty institutions. The Fusion Collaboratory's computational servers will need to balance the equitable sharing of cycles dedicated to offline or retrospective analysis, with the *preemptive* data analysis requirements of operating experiments. Fortunately, computer scientists have developed a promising strategy for dealing with these requirements.

The Globus project has been working for the past 5 years to solve precisely the problems faced by the National Fusion Collaboratory. Globus is facilitating scientific collaboration within flexible "virtual organizations" by connecting globally dispersed collaborators to complex and large-scale data, computing, and visualization resources [14,27]. Components of the Globus project are being utilized, studied, developed, and enhanced at over 80 institutions worldwide.

Globus components provide the capabilities to:

- Create "grids" connecting computational resources with their users.
- Track the capabilities of resources within a grid.
- Specify the resource needs of user's computing tasks.
- Mutually authenticate both users and resources.

Globus is distributed in a modular and open "toolkit" form. This makes it easy to integrate resource allocation services into scientific environments and applications. Globus has been integrated with technologies such as the Condor high-throughput computing environment [31] and the widely used Portable Batch System (PBS) job scheduler [32]. Both of these integrations demonstrate the power and value of the open protocol toolkit approach, and offer significant opportunities for constructing the Fusion Collaboratory.

Each new application of computational grid technology facilitates the deeper understanding of real-world needs and hones the technology to increasing levels of functionality, performance and robustness. The design of the Fusion Collaboratory has identified important research topics of interest to the Globus group, such as resource management for near-real-time experiment steering and policy-based management of resources in multi-organization national collaborations. By bringing together the state of the art in computational plasma physics with the state of the art in computer science, important benefits will accrue to both disciplines.

2.3.2. Strategy

The Fusion Collaboratory is a pilot project emphasizing software development. A major investment in computer hardware is not envisioned. Deployment for testing will be on available hardware such as workstations and small to midrange parallel machines at the computer science and fusion laboratories. The project laboratories will ensure that sufficient hardware resources are available.

The expectation is that success of the Collaboratory will lead to deployment for production use on the DOE Science Grid or on machines at a Topical Center for Fusion Supercomputing towards the end of this three-year pilot project. Given the requirements for pre-emptive scheduling of supercomputers for between-pulse support of fusion experiments, a Topical Center for Fusion, overseen by a consortium of fusion scientists, may be the preferred route.

Globus servers, designed for the use of the fusion community will be set up initially at the computer science laboratories at ANL and LBNL. The fusion laboratories at GA, MIT and PPPL will acquire the necessary expertise to set up their own Globus servers within the first six months of the project. The Globus

servers at ANL and LBNL will be used to alpha test new server software developed by the computer science partners; the fusion laboratories' servers will be used for beta testing.

The general strategy for the deployment of fusion codes will be to work up from small serial codes to small parallel codes to large serial codes to large parallel codes.

2.3.3. Activities of the Computer Science Laboratories

The Fusion Collaboratory will utilize Globus facilities for distributed computing: remote job scheduling, monitoring, exception handling, and accounting. This will enable researchers and their institutions to share the community's computational resources. Collaboration-wide policies to manage the allocations and priorities of users, groups, and research functions will be developed and employed.

The components of the Globus toolkit that can be immediately deployed to create the foundation of the Fusion Collaboratory are:

- Metacomputing Directory Service (MDS), the Globus toolkit's Grid Information Service.
- Grid Security Infrastructure (discussed above).
- Globus Resource Allocation Manager (GRAM).

The *MDS* provides an information infrastructure for computational grids, using the LDAP protocol as a uniform means of querying system information from a rich variety of system components, and for constructing a uniform namespace for resource information across a multiple organizations.

The *Grid Security Infrastructure* is considered in detail above. All resource management functions are executed in the uniform mutually authenticated security context provided by GSI.

The *Globus Resource Allocation Manager* allows compute jobs to be run remotely, providing an API for submitting, monitoring, and terminating jobs. GRAM provides a uniform interface to the various local resource management tools such as PBS, LSF, NQE, LoadLeveler, and Condor. It uses a flexible resource specification language (RSL) to provide a method for exchanging information about resource requirements, job process creation, and job control between all of the components in the Globus resource management architecture. GRAM enables remote monitoring and management of jobs already created, and updates MDS with information regarding the availability of the resources it manages.

The entire Globus toolkit has been ported to and tested on a wide variety of Unix platforms, so running the server components of the toolkit on the various Fusion servers should be straightforward. The issues of Globus job submission from non-UNIX clients have been studied and several straightforward solutions exist. These will be developed and deployed in accordance with the needs of Collaboratory users.

Going beyond the components listed above, research within the Globus project relevant to the Fusion Collaboratory includes:

- End-to-end resource management that can knit together knowledge and quality-of-service control of networks, compute power, and storage resources, in order to automate the complex process of sharing resources among multiple organizations, users, and purposes.
- Advanced reservation or pre-emptive allocation of computational resources for high priority use, e.g. for runtime support of scheduled experiments. Preliminary work has been done on advance reservation with the support of the DOE Cliper project [33]. Work in this project will address challenging issues relating to pre-emption and the policies that govern when pre-emption is supported.

A diagram showing the role of Globus components in the Fusion Collaboratory is shown in Fig. 2.

Our application of grid tools to the Fusion Collaboratory will involve several research and deployment challenges, taking Globus tools to new levels of functionality:

- Expressing and implementing the Collaboratory's global and local policies.
- Managing batch vs. pre-emptive job priorities, and mapping these to the resource control mechanisms offered by different networks, processors, schedulers and devices.
- Providing status display and accountability back to users.
- Monitoring the adherence of resources to policies.

- Allowing resources to be easily added to and removed from the grid.
- Providing advance reservations, including the co-reservation/co-allocation of multiple resources at once.
- Simplifying user account administration using "dynamic accounts" to make multi-site user authorization easier.

These are all important problems of broad interest, and this project provides an excellent opportunity to explore them in a realistic, focused and highly collaborative environment of significant but manageable complexity.

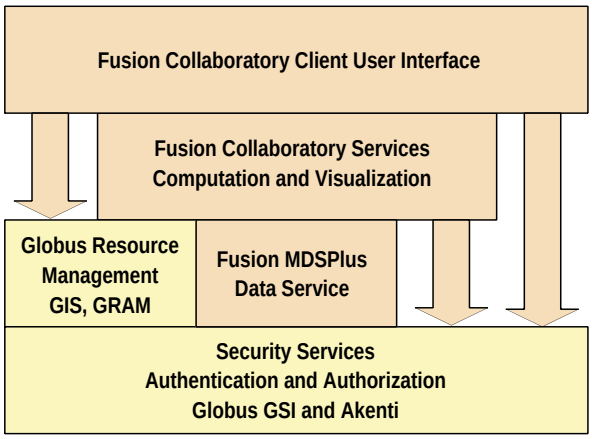

Fig. 2. Collaboratory services are layered, hiding complexity from end users while providing powerful toolkits for application developers.

2.3.4. Activities of Fusion Laboratories

GA, MIT, and PPPL will develop the expertise to administer and operate Globus servers, and will actually deploy fusion codes in their respective Globus servers and on servers at LBNL and ANL, as outlined in the statement of work and schedule. Fusion experimental data will continue to be stored in MDSplus data servers and in SQL Server relational databases. It will be the responsibility of MIT, in close collaboration with the computer science partners, to adapt the fusion program's MDSplus data servers and SQL Servers to meet collaboratory requirements.

The fusion laboratories will deploy specific codes in the Collaboratory. These codes fall into two broad categories: stability codes, which assess the MHD stability of plasmas in given or observed configurations, and transport codes, which assess the transport and quality of confinement of particles, momentum, and energy within the plasma.

The stability codes to be deployed are as follows: DCON [34], GATO [35], PEST-1,2,3 [36,37]. The transport codes to be deployed are: the "phenomenological" transport data analysis codes TRANSP [38] and ONETWO [39], and the theory (micro-turbulence) based transport code GS2 [40]. The phenomenological codes will include plasma heating models such as the RF ion cyclotron resonant heating code TORIC [41], which may also be made available in standalone form.

In addition to actual deployment of the fusion codes, it will be the responsibility of the fusion laboratories to develop client software for users. This means user interfaces that use the tools provided by the computer science laboratories to allow:

- Secure login and user authentication.
- Identification and preparation of data resources.
- Initiation of code run requests (batch or preemptive priority).
- Cancellation of run requests.
- Ability to monitor progress of runs and view intermediate results.
- Ability to halt and delete, or halt and archive where appropriate, runs already started.
- Ability to receive notification of run completion (normal or abnormal).
- Access run results (visualization; post-processing).
- Automation of run production process.

The fusion laboratories will arrange for live user tests, both for offline analysis computation and for between-pulse analysis.

2.4. Scientific Visualization

2.4.1. Introduction

The highly distributed yet collaborative nature of fusion research necessitates a shared visualization environment linking fusion researchers together into a coherent scientific team. The current software environment deployed in the fusion community creates a barrier to effective scientific collaboration and therefore negatively impacts the quality of research. The visualization components of this proposal will strive to eliminate or dramatically reduce the hurdles that presently exist for collaborative visualization (e.g., the lack of flexible and collaborative visualization tools). The project will leverage existing technology where possible, however, because some important requirements cannot be fulfilled by current software capabilities, a component of our proposed work will depend on innovative computer science research.

One main component of our work is to create a collaborative control room for the three main experimental fusion facilities. The collaborative control room concept allows the large on–site group to interactively work with small to large off–site groups. Within the control room, a large high–resolution "display wall" will be used to allow a large group of scientists to explore information in collaboration more effectively than if they were crowded around a single workstation display. The proposed technology will allow off–site research teams to be intimately involved in the decision making process during the important 15 minute between-pulse analysis cycle thereby overcoming a significant current collaborative barrier for fusion scientists. Scientists will be able to simultaneously share and interact with the same visualizations (whether from experiment or simulation) creating a more efficient overall working environment that will lend itself to more rapid scientific discovery.

To reduce costs, testing and proof–of–principle demonstrations to the fusion community will be done by connecting a multiple pairs of existing large–scale, high–resolution display walls. These tests will require the integration of remote display and control technologies with the existing Access Grid (AG) collaboration framework for simultaneous visualization on display systems distributed across the Collaboratory. The goal of this work is to demonstrate the feasibility of deploying large display walls in the control rooms of the three large experimental facilities and at significant remote sites. Upon completion of this feasibility goal, midway through the lifetime of this project, we will seek other funds to purchase the hardware necessary to fully deploy the technology.

Control room collaboration involving between-pulse data analysis is the most time critical environment within fusion research and will therefore place the greatest stress on our emerging collaborative technologies during testing. Successful tests under control room conditions will mean that this technology will also work for less time critical applications such as working group meetings involving, for example, experimental data analysis, simulation code development, or the comparison of theory and experiment.

Within the control room visual display and AG environments, new software tools will be created to allow simultaneous sharing of complex scientific visualizations. Such tools include the ability to assimilate the "big picture" view of between-pulse data analysis that involves hundreds of different experimental signals and signals generated from simulations. The ability to simultaneously visualize and interactively analyze experimental and simulation data sets will require new visualization techniques and software that can quantitatively compare experimental and simulation data. Furthermore, new visualization techniques and methods will be required for visually representing errors and uncertainties in, sometimes complex, experimental and simulation data.

2.4.2. Remote and Collaborative Visualization

The Access Grid [42,43,44] is an ensemble of network, computing and interaction resources that supports group-to-group collaboration and communication. It was originally developed by the Futures Laboratory in the Mathematics and Computer Science Division at ANL with support from the DOE 2000 program, and is now in use at over 40 institutions including major DOE National Laboratories and NSF PACI Universities. The AG consists of large-format multimedia displays, presentation and interactive software environments, interfaces to grid middleware, and interfaces to remote visualization environments. Large-format displays integrated with intelligent or active meeting rooms are a central feature of AG nodes.

While the AG is often used to link small groups of researchers together it is possible to participate from multimedia enabled PC's. The cost of an AG node today (~$40K) will drop significantly as a result of exploiting new commodity media processors such as those found in game consoles therefore encouraging the deployment of AG nodes in offices and small conference rooms.

Work is proposed in two areas especially targeted at advanced requirements of the Fusion Collaboratory. First, "workspace docking" mechanisms will be developed that are aimed at more smoothly integrating the AG large-format display and shared applications environment with the Collaboratory participant's existing desktop computing environment. The workspace docking system will support a "drag and drop" type operation for "publishing" existing desktop applications to the shared display of the AG. Second, we plan to integrate existing AG collaborative tools and services with new tiled-display oriented parallel window management system called Xplit [45]. This will enable the use of scalable tiled-displays for both high-resolution local and remote visualization and collaboration services such as multiple video streams and distributed desktop tools, which currently rely on the shared memory of multiple graphics adapters on AG node display computers.

Workspace Docking. AG users often desire to "share" some portion of their personal workspace (current desktop applications and data) with other AG users, nodes, or sites, both local and remote. This desire is often extemporaneous and any proposed solution must recognize and enable the ad hoc nature of this desire. The most compelling mode of sharing is the ability to share an application merely by dragging the windows from a laptop onto the shared space in the room. As the application is migrated to the shared space, mouse and keyboard control of the application remain with the user who shared the application, and can be passed around or not as desired. A more mundane sharing mode is the notion of a global shared document space and clipboard. Information can be shared by dragging data regions or entire documents and applications to the shared space. Several options will be investigated such as exporting display primitives, automated application architecture modification and automated applications and data migration, with a focus towards techniques that benefit the majority of the fusion applications.

Most existing desktop application sharing systems are based on the simple idea of replicating the display primitives of the server's application on the client display. This involves creating a copy of the pixels (or other display primitives such as OpenGL calls) on the serving machine and transmitting them via the network to one or more clients. Various forms of compression may be used to improve performance but essentially this scheme is limited by network bandwidth and server performance. This simple scheme can in principle provide low to moderate performance shared viewing of arbitrary applications served directly from a user's desktop machine into the AG environment via a relatively simple multicast enable AG client application. The downside of this approach is that it does not permit asynchronous investigations. The performance of the group display may be limited by the computer and networking performance of the desktop server. Furthermore, since the client application is only typically provided pixels there may be color depth and color mapping compatibility problems. Therefore, this type of sharing scheme will be most useful for the occasional ad hoc sharing of an arbitrary application to a limited number of client sites. It should be noted that this architecture has been successfully adopted to provide a remote vtk visualization tool by using the compression and transmission infrastructure of the vic tool [46]. A variation of this approach be can be implemented using a modified WireGL [47] for OpenGL applications.

Access Grid Tools on Tiled Displays. Collaboration technologies are all about improving communication, and often that communication takes place through the sharing of applications and the joint exploration of data. The existing AG tool suite provides multipoint communication for text, video, audio, distributed PowerPoint and a handful of custom visualization and data analysis applications that have adopted an explicit a many–to–one client server architecture. The AG can run any existing collaborative application tool targeted for the standard Microsoft Windows desktop environment, any standalone Windows application, or X-Windows application, and thus can be minimally used to provide a larger format display environment for group viewing of desktop applications. But these existing tools can not take advantage of the AG collaboration framework that provides a common spatial organizing metaphor, common security infrastructure for secure venues, persistence and recording facilities, and venue based navigation and directory services. In addition, there is a need to bring the collaboration tools and application environment of the AG to tiled displays and to couple the AG communications channels to the new types of visualization applica-

tions that will exploit them. The Xplit architecture is being developed to address this requirement. The principal goal of the Xplit architecture is to enable high performance access to clustered display resources while providing an apparent single point of contact for display services. The development work timetable for Xplit proposed herein will be targeted towards rapid deployment of features for integrating tiled displays into the AG. When fully implemented, Xplit will provide a general-purpose, extensible, high-performance API for incorporating tiled displays into a wide range of applications.

The chief architectural feature that enables high performance (i.e. direct access to rendering acceleration hardware) in Xplit is the separation of content delivery path from the service management path. This choice drives the rest of the design and implementation. Two additional key ingredients to the design are: (1) region as the tile-blind token of display real estate, and (2) virtual display server as single point contact. As Xplit is a network-based scheme, and able to support architecturally monolithic displays as large as technology will allow, it is no longer necessary or advisable to imagine that the entire display acreage is attached in any fundamental way to a single processor, process, or machine. These attributes lead directly to a more egalitarian perspective on the relationship of display real estate to application processes. Specifically, Xplit will serve requests for regions of display space that may or may not be visible (this is very consistent with notions built into X). In this view, the region becomes a first class object. It behaves as a virtual display, with window-like properties, possibly instantiated somewhere on the physical display.

The Xplit virtual display server will manage the physical display, driven by a cluster of rendering engines. It alone will understand the relationship between individual pixels and display nodes. It will be implemented as an extension of the standard X server by the addition of a network display device driver, one netX driver for each display node, allowing the server to manage collectively the resources of the entire display farm. The Xplit client interface will encourage the view that the display is monolithic, hiding the multi-node nature of the display farm from the application through the transparent management of region baggage objects. These encapsulate the relationships between the target display region, the display farm nodes, and the overlapping tiles serving the pixels in the region.

Display Wall. Large high-resolution "display wall" devices are emerging for a variety of applications in scientific and collaborative workspaces (Fig. 3) like the control room environments described in Section 2.4.1. A high-resolution display permits the simultaneous viewing of a variety of sources of information, permitting the viewers both to compare data from different sources and to acquire a better overall understanding of the information than is possible by viewing the individual sources sequentially. The goal of our work is to develop the technology infrastructure necessary for bringing a display wall into the control room where it can allow scientists to integrate the many sources of data available in the precious minutes between pulses.

Fig. 3. Tiled high-resolution display walls allow a large group of scientists to explore information in collaboration more effectively.

While a number of high-end commercial solutions exist for building display wall systems, recent research has focused on the use of low-cost commodity components: a cluster of PCs, PC graphics accelerators, consumer video and sound equipment, and portable presentation projectors [48]. This approach has the advantage of low cost and of tracking technology well, as high-volume commodity components typically have better price-performance ratios and improve at faster rates than special-purpose hardware. The challenge with commodity components is to construct a high-quality collaborative environment that delivers display, rendering, input and sound performance competitive with that delivered by the custom-designed high-end machine approach. A further goal of our work is to construct technology that scales better in terms of pixel resolution and display performance than the custom solutions.

Prototype display walls have been constructed at Princeton University Computer Science (PCS), PPPL, and ANL, utilizing between 9 and 24 projectors per tiled display. While current results are promising [49,48], a number of research challenges remain, including human interfaces [48], load-balance parallel rendering [50,51], projector alignment and color balancing for seamless tiling [52]. Perhaps foremost among the challenges in the control room context is the development of protocols for simultaneous streaming of data from a variety of local and remote sources to a single display surface. The division of client-server applications at the ideal layer as well as efficient protocols for integration of raw pixel data, 2D primitives, 3D primitives, and application control flow is the key to addressing this challenge. A critical component of this work is the combination of Access Grid technology with tiled display technology.

Research in this area will proceed in three stages. Initially, the protocol challenge identified above for purely local data sources will be investigated. Next, will be developed the infrastructure necessary to support these collaborative visualizations involving remote sites: (1) a high-speed connection between the PCS and PPPL display walls, and (2) integration of tiled display technology with access grid for connecting multiple sites, e.g. ANL, PCS, PPPL, and others. Finally, based on these demonstrable prototypes, outside funding will be explored for hardware to establish tiled displays in the three main control rooms.

2.4.3. New Visualization Tools

Within the control room display and AG environments, visualization tools for experimental and simulation data need to be created that will represent a significant increase in capability and efficiency for the fusion community. New software tools will be created to simultaneously share complex scientific visualization to all nodes on the AG and will be extended to common desktop displays at sites not located at the major experimental facilities. Specifically, a collaboration visualization tool will be created for experimental operations giving both local and remote research teams an interactive, integrated big picture view of the vast amount of data that is being acquired. Research challenges include the integration of numerous heterogeneous data streams to create a big picture view of the local and remote data landscape. These streams include experimental and simulation data from multiple MDSplus and relational database servers. The new visualization tool will be structured to allow the scientist to examine large quantities of data from many signals. This will require the development and evaluation of techniques that make use of multiple collaborative representations for enhanced multi-participant interpretation of datasets.

Time-critical rendering and computation permit a user-defined tradeoff between accuracy and interactivity. Research is required to define this tradeoff for locally attached graphics hardware and for remote access to workstations, AG nodes, and high resolution wall systems. The current model of remote visualization is to provide one of two methods: render on a visualization server and download images to be displayed on the graphics workstation, or preprocess on a visualization server and download geometry to be rendered locally on the graphics workstation. The initial research plan is to investigate and develop visualization algorithms that progressively refine visual quality to a user-specified level or at a user-specified frame rate, allowing the user to browse the visualization parameter space efficiently and effectively. Further research in remote visualization will leverage the research in multiresolution methods to create progressive refinements of visualizations that adapt to a set of predefined user and computational/networking criteria, such as frame rate, data size, and bandwidth.

The combination of developing comparison capabilities and creating new visual representations of error and uncertainty will be fundamental to comparing experimental data with simulation results thereby allowing validation of new theoretical models and enabling better understanding of data to facilitate better decision making during experimental operations. Software tools will be created to visualize complex simulations of plasma behavior, initially focusing on microturbulence and macrostability simulation work started under PSACI and proposed under SciDAC. These new tools will allow researchers to compare simulation results with experimental data including the measured experimental uncertainties.

With few exceptions, most visualization research has ignored the issues of visual comparison and representation of uncertainty from data. This is partly because of inherent difficulty in defining, characterizing, and controlling comparisons between different data sets and their corresponding uncertainty in the experimental, simulation, and/or visualization process. Another reason is the lack of methods that allow for easy comparison and representation of uncertainty in data for visualizations. Indeed, one lab

scientist stated [53] that the only way he could compare two three-dimensional visualizations was to print the visualizations out on transparencies, lay them on top of each other and hold them up to the light to compare them. He dubbed this the "view graph norm" and noted that even such simple techniques were not present within visualization systems.

Our goal is to provide users with visualizations that incorporate and compare data from multiple experimental and simulation sources and to reflect uncertainty information to aid in data analysis and decision making. We seek to provide visualization capabilities that present a more complete and accurate rendition of data for users to analyze. Sources of uncertainty include uncertainty in acquisition (instrument measurement error, numerical analysis error, statistical variation), uncertainty in transformation (errors introduced from resampling, filtering, quantization, rescaling), and uncertainty in visualization.

Visual techniques will be developed and evaluated for comparing experimental and simulation 2D and 3D data while incorporating statistical, numerical and/or measurement errors [54,55,56]. Initial methods that will be developed for visual comparison include

- Sets of filters to resample and rescale data
- Overlay and compare 2D and 3D data (automating the view-graph norm)
- Highlight certain physically or numerically important features
- Quantitatively interrogate data sets for numerical values and ranges
- Create new visual representations based upon comparison data (subtraction & interpolated maps)

New visual representations for characterizing error and uncertainty will be explored including

- Physically based glyphs [57]
- Modification of data and/or visualization attributes.
- Improving psycho-visual metaphors
- Enhanced visual representation of statistical data

2.5. Statement of Work and Schedule

The proposed project will cover a period of three years and will be composed of three primary tasks as described below. This section describes the SOW for the project as a whole. Although the individual partners will be funded separately, space limitations restrict our ability to list individual SOW's for each organization.

Task 1— Security

Task 1.1 Deployment of Security Infrastructure
- 1.1.1. Provide certificate authority, GSI services (ANL, LBNL - Year 1)
- 1.1.2. Deploy dedicated fusion CA and GSI services (GA, MIT, PPPL - Year 1 and 2)

Task 1.2 Adapt fusion tools to security infrastructure
- 1.2.1. Modify mdsconnect to use GSI (LBNL, MIT - Year 1)
- 1.2.2. Implement Microsoft SQL Server authentication strategy (ANL, LBNL, MIT - Year 1, 2)
- 1.2.3. Implement Akenti authorization services for MDSplus, SQL Server (LBNL, MIT - Year 2, 3)

Task 1.3 Extension of authentication/authorization tools
- 1.3.1. Add Akenti authorization (ANL, LBNL - Year 1)
- 1.3.2. Refine Akenti tools, user interface (LBNL - Year 1, 2 and 3)
- 1.3.3. Extend credential delegation (ANL, LBNL - Year 2 and 3)
- 1.3.4. Port tools to other platforms (ANL, LBNL, MIT - Year 1, 2 and 3)

Task 2— Remote and Distributed Computing

Task 2.1 Deployment of remote computing infrastructure
- 2.1.1. Deploy GRAM, MDS, GridFTP servers (GA, MIT, PPPL, ANL, LBNL, Utah - Year 1)
- 2.1.2. Develop and deploy Fusion Collaboratory MDS (ANL, LBNL - Year 1)
- 2.1.3. Develop and deploy unified GridFTP–based data access mechanism (ANL - Year 1)

Task 2.2 Computation Management
- 2.2.1. Extend GRAM protocal and implementation to support preemption (ANL - Year 1)
- 2.2.2. Extend resource managers used at Fusion sites to support preemption (ANL - Year 1)
- 2.2.3. Development of brokers and computation managers (ANL - Year 1 & 2)
- 2.2.4. Develop "dynamic account" mechanisms (ANL - Year 2)
- 2.2.5. Develop accounting mechanisms for distributed runs (ANL - Year 2 & 3)
- 2.2.6. Develop mechanisms for specification/enforcement of policies (ANL, LBNL-Year 2 & 3)
- 2.2.7. Integrate event delivery mechanisms for notification of run results (ANL - Year 3)
- 2.2.8. Develop "run directory" recording characteristics (ANL - Year 3)

Task 2.3 Deployment of MHD stability codes in the Collaboratory framework
- 2.3.1. DCON and GATO (GA - Year 1)
- 2.3.2. PEST–1,2,3 (PPPL - Year 1)
- 2.3.3. Preemptive scheduling for between-pulse analysis (GA, MIT, PPPL - Year 2)
- 2.3.4. Time parallel DCON and GATO with between-pulse EFIT (GA - Year 2)
- 2.3.5. Time parallel PEST–1,2,3 (PPPL - Year 2)
- 2.3.6. MHD codes using between-pulse transport analysis (GA, PPPL - Year 3)

Task 2.4 Deployment of transport codes in the Collaboratory framework
- 2.4.1. Serial GS2 (PPPL - Year 1)
- 2.4.2. Preemptive scheduling for between-pulse analysis (GA, MIT, PPPL - Year 2)
- 2.4.3. ONETWO (GA - Year 2)
- 2.4.4. Surface parallel GS2 and TRANSP (PPPL - Year 2)
- 2.4.5. Parallel TRANSP (PPPL - Year 3)
- 2.4.6. Parallel ONETWO (GA - Year 3)
- 2.4.7. MHD codes using between-pulse transport analysis (GA, PPPL - Year 3)

Task 3— Scientific Visualization

Task 3.1 Development and deployment of Remote and Collaborative Visualization Capabilities
- 3.1.1. Create and demonstrate prototype interface to AG on tiled displays (ANL, PCS - Year 1)
- 3.1.2. Demonstrate prototype desktop visualization application with AG (ANL, Utah - Year 1)
- 3.1.3. Demonstrate prototype of AG workspace docking (ANL - Year 1)
- 3.1.4. Create and demonstrate a control room wall mock–up (ANL, GA, PCS, PPPL, - Year 2)
- 3.1.5. Demonstrate Xplit architecture for virtual display (ANL, PCS, - Year 2)
- 3.1.6. Set up AGs at 3 main fusion sites – need external funds (ANL, GA, PPPL, Utah - Year 2)
- 3.1.7. Integrate new visualization tools with display wall/AG (ANL, PCS, PPPL, Utah - Year 3)
- 3.1.8. Integration of AG and remote computing tools (ANL, GA, Utah - Year 3)
- 3.1.9. Integrate Globus security with Access Grid (ANL, PPPL - Year 3)
- 3.1.10. Deploy mini AGs to other sites – need external funds (ANL, GA, PPPL, Utah - Year 3)

Task 3.2 Development and deployment of new visualization tools
- 3.2.1. Develop basic quantitative 2D and 3D visualization (Utah - Year 1)
- 3.2.2. Modify existing tools to work in new structure (ANL, GA, PCS, PPPL, Utah - Year 1)
- 3.2.3. Demonstrate workstation based visualization server (ANL, GA PCS,, Utah - Year 1)
- 3.2.4. Create new visual representations based upon comparison data (Utah - Year 2)
- 3.2.5. Create first version of control room "big picture" display (ANL, Utah - Year 2)
- 3.2.6. Create new physically based visual glyphs for uncertainty visualization (Utah - Year 3)
- 3.2.7. Integrate new visualization tools with remote computing & security (ANL, Utah - Year 3)

2.6. Transition Plan to Sustaining Activities

The long–term goal of this proposal is to create a fusion science grid that would be available to the entire community and provide access to all of the important community resources. To facilitate the process of incorporating new users and new services, workshops will be held in which the vision and technologies of the Collaboratory are discussed. Reference and training documentation will be produced including an online directory of services. In concert with major stakeholders, administrative mechanisms will be

developed for establishing and verifying fair use of shared resources including acceptable use policies. To provide access to our entire community, the PKI/CA infrastructure will need to be extended, if possible in concert with the SC-wide CA proposed by ESnet. Additional services will be added to the collaboratory including more production codes, new computational resources and data from the smaller experiments. It is anticipated that any new experiments or new code development groups will participate as they come on line. The project will be particularly interested in working with fusion SCIDAC application developers to bring their high-end simulation codes into the Collaboratory framework. The project will also work with analysis code developers to optimize their tools for between-pulse analysis to take the maximum advantage of the new capabilities. New visualization software and hardware will be deployed to additional sites as needed. An ongoing dialogue will be maintained between the fusion and CS partners. Experience from the collaboratory can be used to drive improvements in the current generation of tools and to prompt further research and development. At the same time, the partners will explore the use of additional technologies that are developed as part of the SCIDAC program.

2.7. Management Plan

2.7.1. Project Management

The three main experimental fusion laboratories (C–Mod, DIII–D, NSTX) have been working together for many years both in the areas of plasma physics and computer science and have developed a close set of working relationships. Similarly, the computer science laboratories (ANL, LBNL, PCS, Utah) have formed a strong working relationship on a wide variety of research programs. Therefore, the principal management challenge in this project is to insure that these two groups work effectively and efficiently together as a unified team.

Schissel will be the overall PI aided by a management committee comprised of the lead Co–PIs from each site. Resource allocation and direction will be the responsibility of the management committee. To create a cohesive team, Schissel will organize and conduct monthly management committee meetings via video-calls/phone-calls that will deal with high–level planning and coordination issues associated with this effort. These calls will provide the continuity needed to keep the project on track at a strategic level. Complimenting these calls will be one face–to–face management committee meeting per year that will bring the entire team together for discussions that will examine any major direction changes and new ideas development. Meeting locations will be chosen at both the end–user facilities and the computer science research centers so as to promote as much cross–fertilization of ideas as possible.

2.7.2. Oversight Committee

An oversight committee has been selected from the fusion community to advise on requirements and provide feedback. Members of this committee have been selected so as to represent the geographically diverse fusion research community. Oversight committee members include T.A. Casper (LLNL), B. Dorland (Maryland), B. Granetz (C–Mod), C.M. Greenfield (DIII–D), S. Kaye (NSTX), M. Murakami (ORNL), and C. Sovinec (LANL). At least once a year the oversight committee will meet with the Management committee to review progress and provide feedback from the fusion user and code developer communities.

2.7.3. Technical Working Groups

Three technical working groups will be organized around our three main research thrusts: Security, Remote and Distributed Computing, Scientific Visualization. These organizations will have the overall responsibility to move the technical and research and development plans forward, and to coordinate the efforts at each site in those areas. The working groups will meet on an as needed basis but not less than once a quarter. These meetings (both electronic and face-to-face) will be augmented by email lists, and document exchange via the National Fusion Collaboratory web site.

3. LITERATURE CITED

[1] President's Committee of Advisors on Science and Technology, "Report to the President on Federal Energy Research and Development for the Challenges of the Twenty–First Century," November (1997).

[2] I. Hutchinson, R. Boivin, F. Bombarda, et al., "Alcator C–Mod," Phys. Plasmas **1** (1994) 1511.

[3] J.L. Luxon, N.H. Brooks, L.G. Davis, R. K. Fisher, F.J. Helton, M.A. Mahdavi, J. Rawls, R.D. Stambaugh, T.S. Taylor, J.C. Wesley, 11th European Conference on Controlled Fusion and Plasma Physics, 5-9 September 1983, Aachen, Federal Republic of Germany (European Physical Society, Aachen, 1983).

[4] S. Kaye, et al., Fusion Technology, Vol. 36, No. 1 pp 16-35 (1999).

[5] National Research Council, "An Assesment of the Deparment of Energy's Office of Fusion Energy Sciences Program," National Academy Press (2000).

[6] R. Fonck, et al., "Remote Operation of the ITER BES Experiment From an Off–Site Location," Rev. Sci. Instrum. **63** (1992).

[7] S. Horne, M. Greenwald, T. Fredian, I. Hutchinson, B. LaBombard, J. Stillerman, Y. Takase, S. Wolfe, T. Casper, D. Butner, W. Meyer, and J. Moller, "Remote Control of Alcator C–Mod from LLNL," Fusion Technol. **32** (1997) 52.

[8] B.B. McHarg, T.A. Casper, S. Davis, D. Greenwood, "Tools for Remote Collaboration on the DIII-D National Fusion Facility." Fusion Engineering and Design **43** (1999) 343.

[9] T. Fredian, J. Stillerman, "MDSplus Remote Collaboration Support–Internet and World Wide Web," Fusion Engineering and Design **43** (1999) 327.

[10] Microsoft Corporation, http://www.microsoft.com

[11] National Energy Research Scientific Computing Center (NERSC) http://www.nersc.gov/.

[12] DOE Science Grid , http://www-itg.lbl.gov/Grid/

[13] I. Foster and C. Kesselman, (eds.), *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann, 1999.

[14] I. Foster and C. Kesselman, Globus: A Toolkit-Based Grid Architecture in *The Grid: Blueprint for a New Computing Infrastructure*, 1999, Morgan Kaufmann, 259-278.

[15] I. Foster, C. Kesselman, G. Tsudik, and S. Tuecke, "A Security Architecture for Computational Grids," Proceedings of the 5th ACM Conference on Computer and Communications Security (1998).

[16] M. Thompson, W. Johnston, S. Mudumbai, G. Hoo, K. Jackson, A. Essiari, "Certificate-Based Access Control for Widely Distributed Resources," Proceedings of the Eight Usenix Security Symposium (1999).

[17] SciDAC 2001 proposals led by ANL "Security and Policy for Group Collaboration" and "A High–Performance Data Grid Toolkit," and led by LBNL "Distributed Security Architectures" and "Reliable and Secure Group Communication."

[18] S. Fitzgerald, I. Foster, C. Kesselman, G. von Laszewski, W. Smith, S. Tuecke, “A Directory Service for Con_guring High-Performance Distributed Computations,” Proc. 6th Symp. on High Performance Distributed Computing, (Institute of Electrical and Electronics Engineers, Inc.) 365-375, (1997).

[19] K. Czajkowski, I. Foster, N. Karonis, C. Kesselman, S. Martin, W. Smith, S. Tuecke, “A Resource Management Architecture for Metacomputing Systems,” Proc. 4th Workshop on Job Scheduling Strategies for Parallel Processing, Springer-Verlag LNCS 1459, 62-82 (1998).

[20] K. Czajkowski, I. Foster, C. Kesselman, “Resource Co-Allocation in Computational Grids,” Proc. 8th IEEE Symp. on High Performance Distributed Computing, (Institute of Electrical and Electronics Engineers, Inc.) (1999).

[21] C. Pancerella, L. Rahn, C. Yang, "The Diesel Combustion Collaboratory: Combustion Researchers Collaborating over the Internet," Proc. of ACM/IEEE SC99 Conference, Portland, Oregon (1999) (this is not actually referenced. I had in mind the lesson that credential delegation was required in any system of distributed computation and data access.)

[22] http://vlt.ucsd.edu/

[23] http://w3.pppl.gov/NTCC

[24] K. Jackson, S. Tuecke, D. Engert, "TLS Delegation Protocol," Internet Draft draft-ggf-tls-delegation-05.txt, http://www.gridforum.org/security/ggf1_2001-03/draft-ggf-x509-tls-delegation-09.pdf

[25] Thompson, Engert, Tuecke, "Internet X.509 Public Key Infrastructure Restricted Delegation Certificate Profile," draft-ggf-x509-res-delegation-01.txt, http://www.gridforum.org/security/ggf1_2001-03/draft-ggf-x509-res-delegation-01.pdf

[26] Tuecke, Engert, Thompson, "Internet X.509 Public Key Infrastructure Impersonation Certificate Profile," draft-ggf-x509-impersonation-06.txt, http://www.gridforum.org/security/ggf1_2001-03/draft-ggf-x509-impersonation-06.pdf

[27] I. Foster, C. Kesselman, and S. Tuecke, "The Anatomy of the Grid: Enabling Scalable Virtual Organizations," International Journal of Supercomputer Applications (2001).

[28] W. Allcock, J. Bester, A. Chervenak, I. Foster, C. Kesselman, V. Nefedova, D. Quesnel, S. Tuecke, “Efficient Data Transport and Replica Management for High-Performance Data-Intensive Computing,” Proc. Mass Storage Conference (2001).

[29] R. Housley, W. Ford, W. Polk, D. Solo, "Internet X.509 Public Key Infrastructure Certificate and CRL Profile," RFC 2459, Draft Update draft-ietf-pkik-new-part1-06.txt (1999).

[30] http://www.iPlanet.com/products/iplanet_certificate/home_2_1_1ad.html

[31] J. Basney and M. Livny, "Deploying a High Throughput Computing Cluster," High Performance Cluster Computing, Rajkumar Buyya, Editor, Vol. 1, Chapter 5, Prentice Hall PTR (1999); also see http://www.cs.wisc.edu/condor/

[32] http://pbs.mrj.com/ and http://pbs.mrj.com/overview.html

[33] I. Foster, A. Roy, V. Sander, “A Quality of Service Architecture that Combines Resource Reservation and Application Adaptation,” Proc. 8th Intl Workshop on Quality of Service (2000.)

[34] A.H. Glasser and M.S. Chance, Bull. Am. Phys. Soc. **42**, 1848 (1997) (abstract dMopP12).

[35] L.C. Bernard, F.J. Helton, and R.W. Moore, Computer Phys. Commun. **21**, 377 (1981).

[36] R.C. Grimm, R.L. Dewar, J. Manickam, J. Comput. Phys. **49** (1983) 94.

[37] A. Pletzer, A. Bondeson, R.L. Dewar, J. Comput. Phys. **115** (1994) 530.

[38] R.J. Hawryluk, "An Empirical Approach to Tokamak Transport," Phys. Plasmas Close to Thermonuclear Conditions, ed. by B. Coppi, et al., (CEC, Brussels, 1980), Vol. 1, p. 19.

[39] W. Pfeiffer, R.H. Davidson, R.L. Miller, R.E. Waltz, General Atomics Report GA-A16178 (1980).

[40] M. Kotschenreuther, et al., Comp. Phys. **88** (1995) 128.

[41] M. Brambilla, "A Full-Wave Code for Ion Cyclotron Waves in Toroidal Plasma," Max Planck Intitute Laboratory Report IPP 5/66 (February, 1996).

[42] http://www.accessgrid.org

[43] L. Childers, T. Disz, R. Olson, M.E. Papka, R. Stevens, T. Udeshi, "Access Grid: Immersive Group-to-Group Collaborative Visualization," Proc. of the 4th International Immersive Projection Technology Workshop (2000).

[44] L. Childers, T. Disz, M. Hereld, R. Hudson, I. Judson, R. Olson, M.E. Papka, J. Paris, and R. Stevens, "ActiveSpaces on the Grid: The Construction of Advanced Visualization and Interaction Environments," Proc. PDC Annual Conference: Simulation and Visualization on the Grid, to be published in Springer-Verlag Lecture Notes in Computer Science (2000).

[45] M. Hereld, I. Judson, R. Stevens, "Xplit: An Architecture for Integrating the Access Grid with Tiled Displays," Proc.of the Access Grid (AG) Technical Retreat, Argonne National Laboratory, Argonne, IL (2001). Seel also http://www.mcs.anl.gov/fl/accessgrid/ag-retreat-2001/ag-retreat-2001.htm

[46] R. Olson and M.E. Papka, "Remote Rendering Using Vtk and Vic," presented at the Hot Topics section of *Visualization 2000*, Salt Lake City, Utah (2000).

[47] http://graphics.stanford.edu/software/wiregl

[48] K. Li, H. Chen, Y. Chen, D.W. Clark, P. Cook, S. Damianakis, G. Essl, A. Finkelstein, T. Funkhouser, A. Klein, Z. Liu, E. Praun, R. Samanta, B. Shedd, J. Singh, G. Tzanetakis, J. Zheng, "Early Experiences and Challenges in Building and Using a Scalable Display Wall System," IEEE Computer Graphics and Applications **20** (2000) 671-680.

[49] T. Funkhouser, K. Li, "Large Format Displays, Guest Editor Introduction to Special Issue on Large Format Displays," IEEE Computer Graphics and Applications **24** (2000) 20-21.

[50] R. Samanta, J. Zheng, T. Funkhouser, K. Li, J. Singh, "Load Balancing for Multi–Projector Rendering Systems," SIGGRAPH/Europgraphics Workshop on Graphics Hardware (1999).

[51] R. Samanta, T. Funkhouser, K. Li, J. Singh, "Hybrid Sort–first and Sort–last Parallel Rendering with a Cluster of PCs," SIGGRAPH/Eurographics Workshop on Graphics Hardware (2000).

[52] Y. Chen, D.W. Clark, A. Finkelstein, T. Housel, K. Li, "Automatic Alignment of High–Resolution Multi–Projector Displays Using an Un–calibrated Camera," IEEE Visualization 2000 Conference, Salt Lake City, Utah (2000).

[53] Private communication.

[54] C.R. Ehlschlaeger, M.F. Goodchild, "Uncertainty in Spatial Data: Defining, Visualzation, and Managing Data Errors," Proc. of GIS/LIS, Phoeniz, Arizona (1994) 246-53.

[55] A.T. Pang, C.M. Wittenbrink, S.K. Lodha, "Approaches to Uncertainty Visualization," The Visual Computer (1998).

[56] J. Abigail, "Visualizing Uncertainty in Isosurfaces," M.S. Thesis, University of California, Santa Cruz (2000).

[57] C. Wittenbrink, A.T. Pang, S.K. Lodha, "TITLE," in the IEEE Transactions on Visualization and Computer Graphics **2** (1996) 266-279.