Development of a State of the Art Object Oriented Analysis Framework

Principal Investigators: Tim Hallman, Torre Wenaus
Proposal term: 10/99 to 9/00 (renewal of 1999 proposal)

Proposal summary

This proposal is for a continuation of the LDRD project "Proposal to Develop A State of the Art Object Oriented Analysis Framework" initiated last year to draw on RHIC software experience for the development of generic analysis tools which can be widely applied within and beyond HENP, with particular focus on distributed computing. The aims of this project are being realized in the development of NOVA, a Networked Object-based enVironment for Analysis.

Analysis frameworks manage the data analysis process from algorithm development through data selection and retrieval, filtering based on physics features, to final results analysis, in interactive and batch production environments. They are typically experiment-specific, developed with scarce resources and focused on local needs rather than reuse. Quality tools developed with reuse in mind are, however, prized by the community. This LDRD project gives us the capability to go beyond narrowly focused tools to provide a generic, reusable framework for distributed analysis designed for wide use.

In its first year to date the project has designed a distributed computing architecture (NOVA) that employs the burgeoning technologies of distributed computing to provide generic distributed OO analysis tools to the HENP community and similarly challenging computing environments. All key components of the architecture have been prototyped or implemented and several have been deployed in production or as demonstration prototypes. Because NOVA developers are STAR and ATLAS collaborators these experiments have served as convenient testbeds, but NOVA itself is a generic development is not tied to a particular experiment.

The project has established a recognized distributed analysis expertise at BNL that has led to a US ATLAS Computing funding proposal in this area and has garnered wider interest in the community.

In the second year we propose to complete the implementation guided by application and prototyping in STAR, ATLAS and potentially other environments, and establish NOVA as an experiment-independent tool set for distributed analysis.

Looking at things less seriously one can note that HENP experiments of yesterday were distributed and the computing was centralized. Our project is betting that tomorrow experiments will be centralized and computing will be distributed.

Extended proposal

1. Introduction

The size and geographical distribution of current and next generation HENP experiments and their computing infrastructure makes support for distributed computing capability within the analysis software essential. At the same time, the rapid growth in network capacities and technologies now underway make a powerful distributed computing capability both possible and practical to a degree not seen in previous generations of HEP experiments. The NOVA project is building a tool set for the realization of distributed computing in the Object Oriented analysis environments that all forthcoming experiments have adopted. It will enable use of an experiment's analysis framework interactively or in batch as a thin local client providing transparent access to data and computing resources at a remote center. While an important project goal is early integration and prototyping of NOVA components into existing experiments, its design and implementation is as a set of generic components applicable to the OO analysis environment of any experiment. It incorporates the latest standards and technologies in component software, component middleware and distributed computing to support a new level of distributed data-intensive object oriented physics analysis.

The success of this effort will constitute a first step in establishing an important new computational science component in the research effort at BNL. It will lead to a BNL-supported software product that will provide new capabilities serving BNL and HENP community physicists participating both in BNL-hosted research such as the RHIC program and in worldwide collaborations such as the LHC. It will improve the depth and visibility of the Laboratory's contribution to HENP community software and better position the Laboratory for important roles in computing and software for near and long term projects.

BNL is responsible for much of the core software development activity of the RHIC experiments. For the STAR experiment, overall computing management and the core software infrastructure development team are located at BNL. The BNL team is responsible for the development of the STAR analysis framework now successfully deployed and in use for STAR production analysis. This experience base is being employed in the development of NOVA and the STAR analysis environment is being used as one of the testbeds for NOVA.

NOVA components developed in the first year of the project have been deployed and tested in the STAR environment, and subsequently developed further and deployed outside STAR, in particular as demonstration prototypes in the LHC experiment ATLAS.

BNL is a major participant in the US ATLAS program and is the site of the US ATLAS Project Office. In the last year BNL has been selected as the site of the US Regional Center for ATLAS computing and is on a trajectory to assume a major role in ATLAS computing and software. A key goal of this LDRD and the NOVA project is to bring to

the ATLAS computing program an existing BNL capability in distributed analysis software that can help leverage a BNL role in ATLAS analysis software. The analysis framework has been identified by ATLAS and US ATLAS as an area in which US institutes can play a major role.

2. Architecture and Components

Many well-developed experiments already have established object-oriented frameworks in production or under development. However, the present generation of object oriented analysis frameworks are limited in several respects:

- Support for distributed analysis in large, geographically dispersed collaborations
- Management of the complete analysis process from analysis software development through data set selection, data retrieval, bulk data filtering, analysis production, results analysis and iteration of the process
- Integrated support for tracking, validating, and debugging physics analysis in a collaborative environment in which effective communication and documentation of the analysis process is both important and difficult

The NOVA project does not reinvent or evolve existing analysis frameworks, but rather provides new capabilities in these areas via modular components providing application-neutral interfaces that can be used in isolation to extend the capability of existing analysis systems. NOVA software components are focused on four functional areas:

1) Distributed analysis
A mobile analysis client served by a central analysis server and associated monitoring and control tools for control and monitoring of physics analysis in a distributed, heterogeneous computing environment

2) Distributed software management
Integration of software distribution, management and version coordination tools with a mobile analysis client, problem reporting system, code navigation system, and discussion system to serve a widely distributed community

3) Event data and distributed data access
Central data and file catalogues with associated tools for controlling and monitoring data locality, data browsing, and evolving data models (schema evolution) for data-intensive analysis

4) Software robustness and reusability
Dynamic customization of a stable core analysis framework through shared libraries, with management of 'software signatures' for assured reproducibility and invariance of analysis environment.

The NOVA architecture (Figure 1) consists of independent interoperable components designed for flexibility and ease of reuse. We focus principally on supporting C++ based

analysis since C++ is the analysis software language for all RHIC and LHC experiments and most other large experiments. NOVA is being developed using an iterative process driven by user participation and closely coupled to prototyping in real-world experiments (STAR, ATLAS). NOVA architecture components are distributed among four principal domains:

1) **Data management domain.** Data and file catalog databases and associated interfaces
2) **Analysis server domain.** Centralized analysis production and software management
3) **Mobile analysis domain.** Analysis client and associated tools for distributed analysis
4) **Web middleware domain.** Communication and data exchange infrastructure

Existing experience and the evolutionary path of HENP computing have guided the requirements to be met by tools and technologies employed within NOVA components:

- Should be free or nearly so, such that the buy-in cost of using the system is very small
- Should be widely used, true or defacto standards, with good support and showing good growth
- Should be known within or on a growth path within the HENP community

Following these requirements we have adopted the following tools for application in the development of NOVA:

MySQL - Open Software relational database
Data catalogue, event store navigation, mobile analysis client state persistency, software signature management

XML - Data description/storage tool for text-based data
Software source distribution between mobile analysis client and analysis server. Analysis filter, query specification

Apache - Open Software web server
Distributed client/server communication between mobile analysis client and analysis server. Web-based control and monitoring

CORBA - Distributed inter-process communication
Component middleware for low-volume control data

ROOT - Object oriented toolkit for HENP analysis
Baseline implementation layer for analysis server framework, client analysis tools and data storage (optional, experiment-dependent component)

3. First Year Progress

In its first year the project has succeeded in

- Designing a distributed computing architecture and toolset that has captured interest and/or applications in RHIC and LHC experiments including STAR, ATLAS and LHCb. NOVA has found production applications in STAR and is the basis for

ATLAS demonstration prototypes. The computing leadership of the LHCb experiment is visiting BNL this month to learn about NOVA.

- Implementing distributed computing and analysis tools in the areas of configuration and parameter databases, file catalogs and event meta-data databases, associated web-based browsers and management tools, and communication and distributed development tools
- Deploying all these tools in production or prototype form within STAR, and deploying parameter database tools in prototype form in ATLAS
- Presenting the architecture and toolset to the community. Presentations have been made to STAR and US ATLAS with presentations to CERN ATLAS and LHCb in the next month. Presentation to the whole community at the Computing in High Energy Physics conference in January is planned.
- Leveraging the project to define a BNL role in US ATLAS core software development, now being negotiated with US ATLAS computing management.

3.1 Current Status of Project Domains and Components

Bold items are project component deliverables, either fully implemented in the project or third party tools customized or extended for NOVA. Non-bold items are third party tools used by (or with provision for use by) by NOVA. Italic items are application components used with NOVA.

Data Management Domain
Catalog interface (implemented; deployed in STAR, prototyped in ATLAS)
Data Catalog (implemented; deployed in production in STAR)
Data Repository (interfaced to STAR; partial interface to ATLAS)
Grand Challenge Architecture (third party tool; interfacing this fall)

Analysis Server Domain
Analysis daemon (prototyped)
Analysis catalog (prototyped)
Offline control framework (interfaced to STAR)
Dynamically loaded applications (interfaced to STAR)
CVS code repository (interfaced to STAR)

Mobile Analysis Domain
Mobile analysis client (prototyped)
ROOT analysis (interfaced to STAR)
NanoDST (interfaced to STAR)
Grand Challenge Architecture (third party tool; interfacing this fall)
Web browser (web based browsing implemented)

Web Middleware Domain
Analysis client state server (prototyped)
Client state database (prototyped)

Analysis monitoring module (prototyped)
Bug reporting and tracking system (third party tool, integrated)
HyperNews discussion system (third party tool, integrated)
Web server (third party tool, integrated)

3.2 NOVA Solutions

NOVA work to date has provided solutions for several use-case scenarios typical in HENP computing:

Dynamic Brokerage
Research software is inherently dynamic as users seek to incorporate new elements as ideas develop. The typical outcome would be that a user modifies the structure of objects in his application, and the application stores new objects in the database. Remote applications unaware of new functionality may request objects in an old format. The solution to this problem employed in NOVA is the dynamic object request broker. In this approach applications provide along with the metadata request (name, time, selectors...) also the description (dictionary entry) describing the expected application dataObject. The database server supplies along with the requested dataObject also the description for the database DataObject. The NOVA component – Object Request Broker module - converts the database DataObject according to the application dictionary. Figure 2 shows an example of dynamic brokerage use, building the full GEANT geometry of the ATLAS experiment from the parameters stored in the prototype database developed in NOVA. There are several benefits of this approach:

- Separation of database and analysis applications
- Robust interface (via built-in type checking)
- Dictionary built from simple, standard C header files or IDL files
- Database access is independent of application code version: users can read new dataObjects with an old executable

Static Brokerage
In the case of stable applications, such as the web-based database, NOVA implements the simpler solution of static object request brokerage. In this case only the database server side is dynamic and just the application ID is enough to fulfil the task of, e.g., navigating the current database hierarchy. The NOVA Object Request Broker at the central site serves dynamic HTML DataObjects in format tailored according to web browser ID: Netscape or MS Internet Explorer. Figure 3 shows an example of static brokerage use – navigating the hierarchy of the parameter database developed in NOVA.

NOVA Data Model
The multi-layer architecture of the NOVA database implementation is shown in Figure 4. The MySQL database server provides the physical database implementation in NOVA. The object request broker layer provides database views to the applications in terms of simple objects (C structures). The logical database design providing the capability of storing objects in the relational database is shown on Figure 5.

3.3 Remaining goals for FY99

We are on target with respect to our year one schedule, currently in the implementation and testing phase culminating in a status report and year two plan this month, and delivery of components, documentation and final report next month.

Based on progress to date we expect to complete our year one goals by the end of the project's first year. All components should be implemented and deployed in a real user environment (either STAR or ATLAS or both).

4. Year two program

The first year has focused on design, implementation, and early deployment and testing in environments in which the developers are directly involved (STAR and ATLAS). In the second year we will further develop and refine the functionality and implementation of NOVA components based on experience in these environments and complete the application-neutral interfaces that will enable NOVA deployment in environments in which the developers are not directly involved. In the second year we propose to

- extend the deployment and testing of NOVA components in STAR and ATLAS including scalability testing with management of multi-terabyte data sets and large scale analysis,
- refine component functionality and implementations based on experience from deployment and testing,
- continue wider discussions and presentations in the community and seek other deployment/prototyping environments,
- complete the application-neutral interfaces of NOVA components to establish NOVA as an experiment-independent tool set for distributed analysis.

5. Project effort and budget

The level of effort and budget profile for the second year is similar to the first year. The dedicated manpower devoted to this project consists of one full-time person for one half year (Sasha Vanyashin, a visitor to BNL from Royal Institute of Technology, Stockholm), augmented by contributions from existing computing professionals in the BNL STAR computing group (Jeff Porter, Valery Fine). BNL physicists provide project leadership (Torre Wenaus) and further development effort (Torre Wenaus, Pavel Nevski). Small funding requirements for space, travel and a computer workstation complete the budget.

6. Potential future funding

In discussions with US ATLAS Computing management we have reached agreement on the value of leveraging NOVA and the expertise behind it for ATLAS Computing and we are negotiating the near and long term role and funding appropriate to making this happen. To this end we have requested funding at a 1FTE level from US ATLAS Computing Project DOE funding for the initiation of ATLAS distributed control software development at BNL and the application of NOVA within ATLAS.

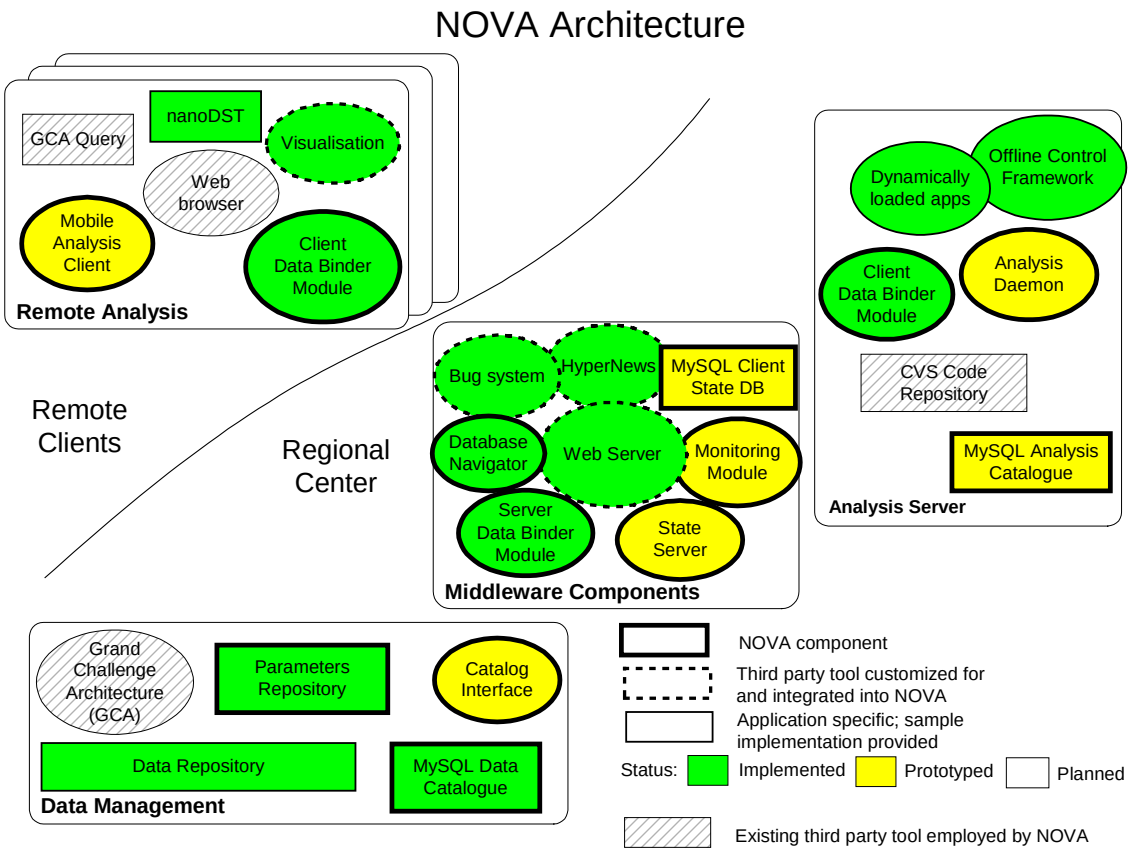

Figure 1: The NOVA Architecture

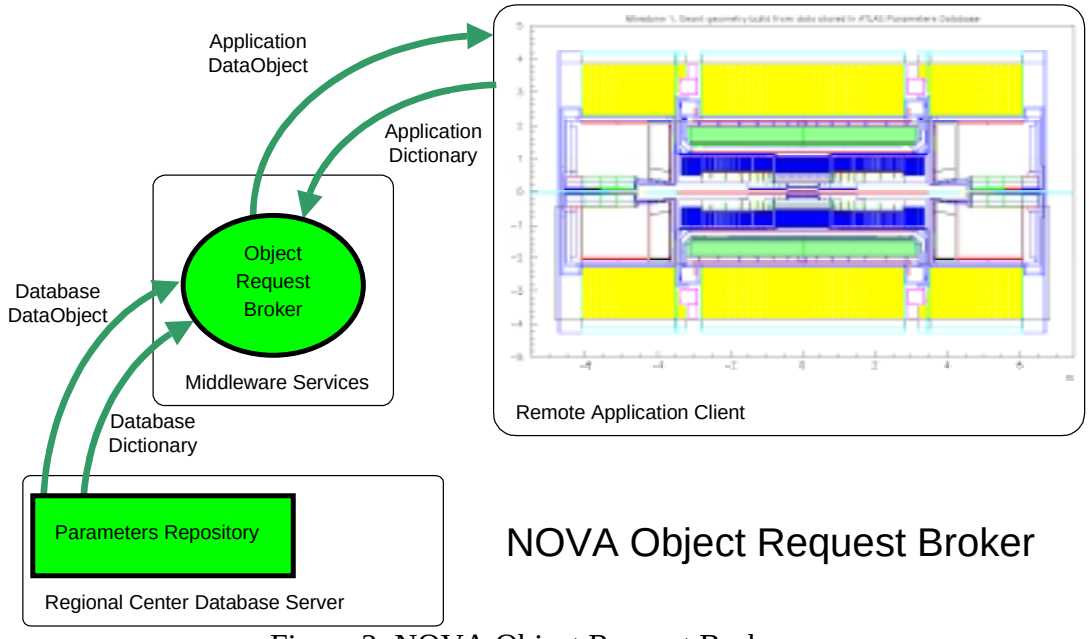

Figure 2: NOVA Object Request Broker

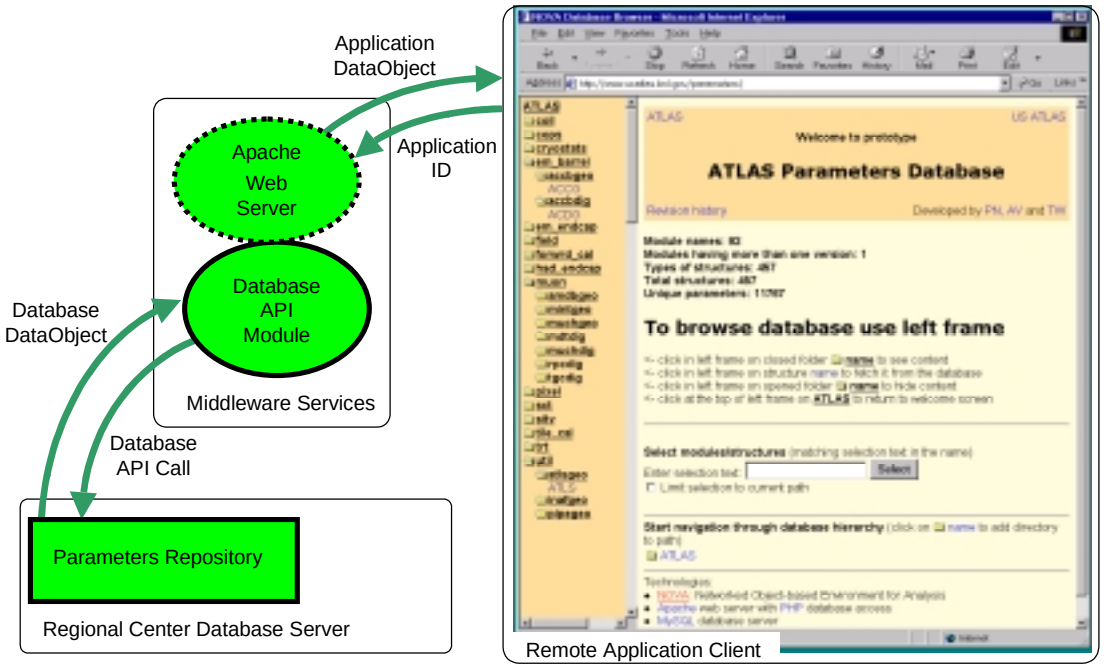

Figure 3: Navigating the parameter database

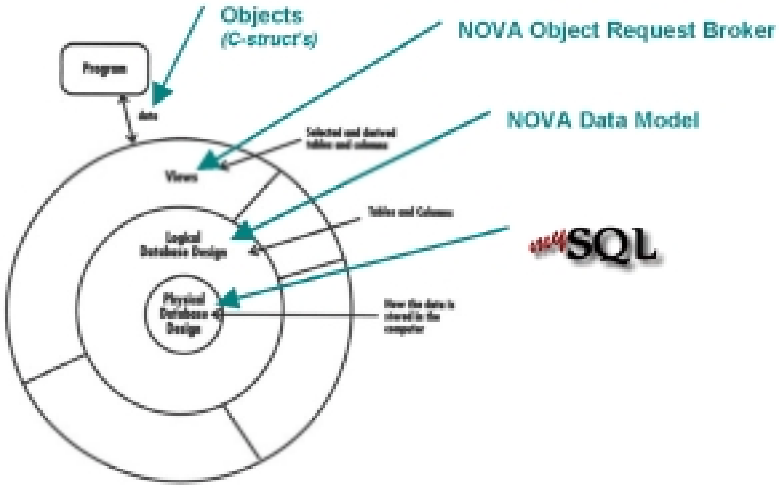

Figure 4: Multi-layer architecture of NOVA database access

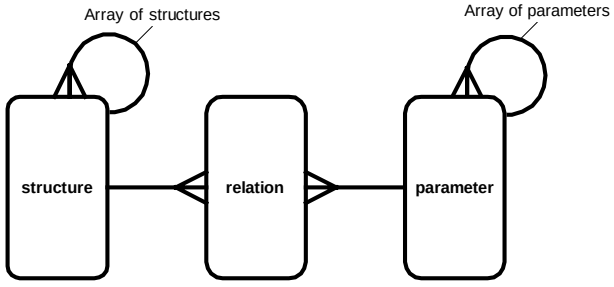

Figure 5: NOVA data model for storing objects in relational database