APOGEE : A Physics-Optimized Grid Environment for Experiments

Introduction

The major particle and nuclear physics experiments of the next twenty years at SLAC, FNAL, BNL, JLAB and the LHC at CERN, will break new ground in our understanding of the fundamental interactions and symmetries governing the nature of matter and spacetime. The physics programs of the major HENP experiments, representing a multi-billion dollar investment by DOE and NSF, require the coordination of tens to hundreds of physics groups, and the efforts of thousands of scientists from every region of the globe. Realizing the scientific wealth of these experiments presents new problems in data access, processing and distribution, and collaboration across national and international networks, on a scale unprecedented in the history of science. Meeting the scientific goals of these experiments requires that numerous challenges in information technology be met. Among these challenges are:

- Providing rapid access to event samples and subsets drawn from massive data stores, rising from 100s of TeraBytes in 2000 to PetaBytes by 2005, to 100 PetaBytes by 2010.
- Transparent access to computing resources, throughout the US, and throughout the world.
- The extraction of small or subtle new physics signals from large and potentially overwhelming backgrounds.
- Enabling access to the data, and to the rest of the physics community across an ensemble of networks of varying capability and reliability, using heterogeneous computing resources.

Finding a solution to these problems is critical for the involvement of young scientists at universities in the ongoing process of search and discovery.

We propose to design, instrument, model, optimize and prototype a data-intensive, hierarchical Grid, composed of national and regional centers linked to the principal center at the experimental site and to local computing resources at the universities. The creation of a hierarchical Grid of data analysis centers linked by Gbps next-generation networks, (complemented, in the case of the LHC Experiments, by a center serving US physicists at CERN), will allow US physicists to play key roles in all stages of the data analysis, from development of the reconstruction programs and software infrastructure, to the extraction of first physics results. This capability will be extended to physicists at their home institutions through the use of an integrated distributed Grid system architecture designed for efficient PetaByte-scale data access and analysis.

While the proposed Grid is targeted at the LHC program, the network-monitoring and control software, and the modeling and optimization systems and methods, will be widely applicable functional components that could drive the design and implementation of future distributed systems in many fields of science, engineering and industry. Data Handling and Computing needs for our field, and other scientific fields, are now of such a size and cost that an engineering and modeling approach to their construction, optimization, and use will have to taken.

The HENP (and LHC) "Data Problem"

High energy and nuclear physics experiments have been coping with increasingly large and complex data volumes, and have been at the forefront in exploiting networks to inter-link their scientific collaborations for the last 15 years. This trend has been driven by the increasing complexity and speed of detectors, data acquisition and transmission systems, all designed to match (and sometimes exceed) the current leading edge capability to process, analyze and manage data. In spite of the continued exponential growth in speed of computing systems and networks, meeting the challenges of the data analysis for the upcoming generations of HENP experiments will require a phase transition — to network-distributed system architectures, data management, and worldwide scientific collaboration on a scale never before encountered.

The crux of the problem is PetaByte-scale data. While proposed computing and data handling systems and networks for future experiments are large by present-day standards, they will not support on-demand access, transport or reconstruction of more than a minute fraction of the data, even in compact pre-processed forms. Studies [Ref] have shown that a distributed hierarchy of facilities linked by high speed networks, is the most promising solution, representing a balance between

- Proximity of the data to central computing and data handling resources, especially for large-scale production-oriented jobs,
- Proximity of frequently accessed data to the users, where it may be processed in desktops, local facilities serving a single physics group, or in "regional centers" depending on the data volume and computational power required for a given job, as well as its relative priority at the central site,
- Making efficient use of limited network bandwidth; especially transoceanic,
- Making appropriate use of regional and local computing and data-handling resources
- Involving scientists and students in each world region, in the analysis and the physics

An effective data analysis by a major collaboration requires a global strategy to handle, prioritize and route requests. The aim of the strategy is to ensure the resources are utilized efficiently while delivering acceptable turnaround times to the entire ensemble of tasks. Once computing and data handling systems and networks of sufficient capability are installed, a successful strategy may be developed based on matching and co-scheduling the requests for data access, processing and transport with the required resources.

Developing a successful global strategy involves the exploration, and determination of an "optimal" solution to several problems:

- Meeting the demands of hundreds of users who need transparent access to local and remote data in disk caches and tape stores;
- Prioritizing hundreds to thousands of requests from the local and remote communities;

- Structuring and organizing the data, and providing the tools for locating, moving and scheduling data transport between tape and disk and across networks, so that it can be accessed and/or delivered rapidly and transparently;
- Ensuring that the overall system is dimensioned correctly to meet the aggregate need.

The solutions and strategies developed will have broad applicability, for the efficient utilization and management of data-intensive distributed systems in research and industry, and eventually in commerce. Systems of similar scale and scope do not exist now, but are expected to be emerging among corporations by the first years of LHC operation.

APOGEE : A Coordinated Approach to the Design and Optimization of a Data Analysis Grid for HENP Experiments

The APOGEE Collaboration, composed of physicists, leaders of computing in high energy and nuclear physics for the major US programs, and computer scientists proposes to design, simulate, prototype and test a computational Data Grid to meet these challenges. The members of our collaboration have a vast store of accumulated experience and leadership in the R&D of Grid components and services, as well as the development of the "regional center" model for future HENP computing. This work has led to the concept of a "computational grid hierarchy" composed of five levels of resources[1]:

- Tier 0: The central facility at the experimental site (CERN in the case of the LHC experiments)
- Tier 1: The national regional centers (such as the US center at FNAL for US CMS and at BNL for US ATLAS)
- Tier 2: A center in one region of a large country (e.g. the US centers proposed by the GriPhyN collaboration)
- Tier 3: The computing resources of a single university group
- Tier 4: An individual workstation.

The five-tiered Grid hierarchy extends the traditional client-server model of at most three tiers (enterprise/workgroup/individual workstation) to a regional, national and worldwide distributed system.

The APOGEE program will focus initially on four main areas:

- ***Simulation of a Computational Grid Hierarchy***, as is being planned for the LHC, since such a Grid is a promising solution for future generations of HENP experiments. The simulation system will allow the LHC and other future experiments to test their data analysis and computing resource-usage strategies, and to reliably estimate the baseline computing and networking resources required, incorporating technology and cost projections over future years.

---

[1] Here we use the terminology of the GriPhyN (Grid Physics Network) Collaboration that is focusing its efforts on the development and implementation of the Tier2 regional centers. See http://www.phys.ufl.edu/~avery/mre/white_paper.html by P. Avery et al.

- ***Instrumentation of the Grid Prototype,*** in order to track the state of the facilities and networks making up the Grid, to test and verify the simulation, and to characterize the performance of the system components under load. Instrumentation of the Grid is needed to identify and log key "events" (or bottlenecks) in the distributed system, and to ensure that the main performance characteristics and saturation points of the facilities and networks are adequately characterized . The Grid instrumentation will be crucial to isolate key factors causing hang-ups or other throughput-limitations in the distributed system.
- ***Using Optimization Tools,*** to find the optimum strategies for data caching, clustering and re-clustering data for efficient delivery, and for the sharing and prioritization in real-time of data analysis-related activities, concurrent with other activities loading the Grid networks. For large HENP experiments, this optimization includes a dynamic balance amongst:- the ongoing production-oriented reconstruction of the full dataset, group-oriented reconstruction of selected data for analysis aimed at specific physics topics, and individual demands for re-reconstruction and analysis by hundreds of individuals situated throughout the Grid
- ***Coordination of US HENP Distributed Data Management Projects.*** Because of its central design and optimization role, standing between the small scale studies and prototype projects and the full-scale implementation of a production-scale Grid for major experiments, APOGEE is the natural place to coordinate a development program focussed on meeting the medium and long-term data management needs of HENP. This project's broader role in the design and development of data-intensive Grids also makes it an appropriate vehicle for HENP, as a field, to establish relationships with demanding applications in other science and engineering fields, and with computer scientists.

APOGEE Deliverables

By using the first year PPDG components and services for distributed data management, and the MONARC simulation tool-set as a starting point, APOGEE will deliver:

- A modular simulation system able to represent the computing facilities, networks, data analysis strategies and network-usage profiles for major HENP experiments.
- A full-system simulation of the Computing Models of CMS and ATLAS, with an emphasis on the Regional Centers and the data analysis strategies that will be used by US physicists in each experiment.
- A set of tools for monitoring, tracking and predicting the state of the facilities at each site, and of the state, loading and response-times of the network links in the distributed system.
- A toolset for optimizing distributed systems, based on new techniques such as genetic algorithms [Ref1], genetic programming [Ref2], or other evolutionary optimization methods. These relatively new methods have proven effective in finding solutions or optimizing complex systems in mathematics, in other fields of science such as geophysics, in engineering applications such as stealth design, in factory-floor operations design, and in code breaking [Ref3]. These methods have not yet been applied widely in the HENP community, although they have been demonstrated to be successful in problems analogous to network-distributed HENP data analysis.

- A set of evaluation criteria, or "cost functions", that are effective in evaluating the Grid hierarchy, or other distributed systems of a similar type. These functions will be based on maximizing throughput of the overall system, balanced against the avoidance of long turnaround times for individual analysis-related tasks. The need for human policy-based criteria to also affect cost and performance will not be forgotten.
- A set of strategies for data organization, task prioritization, and data access and delivery, based on studies of LHC and other experiments' data analysis problems. These will be expressed in terms of guidelines for the design and implementation of an effective data analysis.
- Guidelines for further studies and measurements - to be performed using an actual production-scale Data Grid.

Benefits to the HENP Program of Developing a Nationwide Hierarchical Grid

By siting a center in each region of the US, the number of intellectual focal points for student-faculty interaction and mentoring will be significantly increased. This aspect has been shown to be particularly effective in cases where day trips are possible, and the choice of sites for the centers will take this (as well as several other factors) into account. In order to enhance this effect, we intend to build next-generation remote collaborative systems (now under development). The intense group-oriented activities surrounding the mock data challenge analysis work will help us optimize the systems, and the methods used for remote collaboration, prior to the arrival of real data.

A Nationwide Hierarchical Grid represents a highly leveraged set of human, computing, and networking resources at Universities and Labs, funded through DOE, NSF and possibly other funding agencies. Our HENP funding will go further and be used more effectively if such a work environment can indeed be built, with appropriate coordination between projects and funding agencies.

The Ground Breaking Nature of PPDG, APOGEE and Data-Grid Projects beyond HENP

The ground-breaking nature of the DOE-NGI funded Particle Physics Data Grid (PPDG), our APOGEE project, and related projects in the earth sciences (Earth Science Grid), climatology (Climate Grid), NASA Information Power Grid (IPG), and earthquake engineering, lies in the data-intensive nature of these projects. In computing-intensive grids, the problem is largely confined to locating adequate computational resources, where the data may be processed in place. Once the problem is data-intensive, as encountered in the projects mentioned above, the movement of the data to the computing resources, and delivery of the processed data across networks of limited capacity (relative to the very large data volumes), becomes problematic. The problem of data transport and overall system complexity reaches a new level, never before attacked, when the system involves multiple data sources and sinks, and (even more problematic) hundreds to thousands of human "actors" interacting with the system.

In these systems, the optimum means to carry out each transaction may be a strategic decision, determined according to the "cost" (delay and resource utilization) of locating, acquiring, processing and delivering the data from one or more sources, over the available network paths. In some cases, the alternative of re-computing results to avoid data transport must be considered. The strategy of "virtual" or "on-demand" data must also be considered, namely that of not computing the data at all until triggered to do so by a specific access request. Finding the optimal paths and strategies for data access and delivery and data computation, where accessed and processed (or partially processed) data moves over complex paths through the system, has strong analogies with optimization of industrial processes on a factory floor. Factoring in the realtime demands of the physicists for data and/or processed results, the analogy is closer to more complex applications, such as simulated battlefields and global environmental tracking systems.

The APOGEE Team: A Collaboration Between HENP and Computer Science

The members of our collaboration have worked together with members of the HENP and computer science communities and industry, on a number of R&D projects aimed at high speed multi-TeraByte to PetaByte data access and analysis across multiple storage media and networks. Some of these projects are listed below:

- ***RD45, GIOD (Caltech):*** Investigations, tests and prototyping of wide area network-distributed object databases and mass storage systems for the next generation of HEP experiments, with transparent data access through the object database management system.
- ***MONARC:*** Modeling and optimizing the site and network architectures of the regional center models (hierarchical grids) for the LHC experiments. Performing simulations to develop effective operational strategies for data analysis. Delivering a simulation toolset for future studies by the LHC and other major experiments.
- ***OOFS (SLAC):*** The Objectivity Open File System. Provides a standard interface between an Objectivity database and a variety of storage systems, from a UNIX file system to HPSS, and is well-adapted to systems aimed at distributed data caching, such as Globus (ANL) and the Storage Request Broker (SDSC). OOFS is a vital part of the production data analysis systems for BaBar at SLAC.
- ***CHINA CLIPPER (ANL, LBNL, SLAC):*** Technologies for widely distributed data-intensive applications, leveraging existing Grid (Globus) and other (DPSS) technologies.
- ***HENP Grand Challenge; STACS (ANL, LBNL):*** Distributed TeraByte-scale data management, including query estimation and monitoring, and cache management. The STACS Storage Access and Control System is an integral part of the production data analysis systems for STAR and PHENIX at BNL.
- ***SAM (FNAL):*** The data access framework for PetaByte-scale data access at Fermilab Run2, including metadata cataloging, global resource and cache management, and workflow management and prioritization.
- ***NILE (Cornell):*** System for distributed computing and the management of networked pools of CPU resources, used by the CLEO collaboration.

More recently, as the need for integrated distributed systems has increased, we have formed a strong collaboration with the originators of Grid and related technologies, notably Foster (ANL), Livny (Wisconsin/CS) and Moore (SDSC). This has led to the DOE/NGI-funded Particle Physics Data Grid (PPDG), a joint project between HENP and computer science, to build and deploy a network and middleware infrastructure to support the data access, processing and analysis needs of upcoming and future generations of HENP experiments. The principal PPDG deliverables, including (1) a high speed bulk data transport service constructed on a network and resource-state aware instrumented foundation, (2) multi-site cached file access service, will use a broad range of middleware components drawn from the projects listed above, and from the Condor (Wisconsin), Globus (ANL) and SRB (SDSC) systems developed by the computer scientists in our collaboration.

The integration of HENP and computer science efforts is illustrated by the plan for the use of components in the first-year PPDG system (from the PPDG proposal):

Object-based and File-based Application Services	**Objectivity DB (plus SLAC/BaBar enhancements); Grand Challenge Query Object, Event Iterator, Query Monitor; Fermilab SAM system**
File Access Service	**Components of OOFS (SLAC)**
Cache Manager	**Grand Challenge Cache Manager (LBNL)**
Mass-Storage Manager	**HPSS, Enstore, OSM (Site-dependent)**
Matchmaking Service	**Condor (U. Wisconsin)**
File Replication Index	**MCAT (SDSC)**
Transfer Cost Estimation Service	**Globus (ANL)**
File Fetching Service	**Components of OOFS (SLAC)**
File Mover(s)	**SRB, Storage Resource Broker, (SDSC); Site-specific movers**
End-to-End Network Services	**Globus tools for end-to-end QoS reservation**
Security and Authentication	**Globus (ANL)**

The APOGEE project's main focus is the design, simulation, instrumentation, and optimization of the Grid for the LHC (and potentially other) major experiments. Beyond this restricted focus, the project's studies and developments will be extended to both learn from and benefit Babar, Run2 (D0, CDF), RHIC (STAR, PHENIX) and JLAB, particularly as they each accumulate data reaching close to and above the PetaByte range in future years. The APOGEE work will go beyond the high energy and nuclear physics communities in preparing for the LIGO data analysis in the context of GriPhyN, and in constructing efficiently accessible astrophysics databases in the context of the ALDAP[2] project.

[2] Accessing Large Distributed Archives for Astronomy and Particle Physics. This NSF/KDI-funded project involving Johns Hopkins, Caltech and the FNAL Sloan Digital Sky Survey, is focusing on the means to efficiently structure, access and dynamically cluster data in distributed object and object-relational databases, for both astrophysics and particle physics.

APOGEE also will provide a forum for forward-looking ideas on the construction, and operation of "self-aware" distributed systems. Some the current ideas:- using organized ensembles of "autonomous agents" to assist with the state tracking, workflow management, and dynamic adaptation to changing site and network states, have arisen and been developed in a broad range of fields of science, engineering and military applications. Use of these ideas to build distributed systems, to deal with problems of a new level of complexity and scale, is as important to the geophysical and earth sciences, and network and other fields of engineering, as it is to high energy and nuclear physics, and accelerator technology. Without a forum covering and extending beyond the limits of the HENP community, it will not be possible to exploit these newly conceived system concepts to develop cost-effective solutions to HENP's major challenges. Without a determined effort to become players in this new arena, HENP computing systems will remain mere isolated scaled-up versions of current systems, with little pooling of hardware, ideas, development and operational costs. This will severely limit their functionality and our science discovery capabilities.

APOGEE in Relation to Other Projects Aimed at the LHC and Other Distributed Computing Problems

This proposal builds on the initial work on a data-intensive Grid testbed for HENP applications by the DOE/NGI Particle Physics Data Grid (PPDG) project, and on the early simulation work applied to the LHC Computing Models by the MONARC project at CERN. It will also utilize, directly or indirectly, the results of the GIOD, China Clipper, Nile, SAM, HENP Grand Challenge, and STACS projects, and components and services from Globus, Condor, SRB and OOFS (summarized above) for distributed data processing, transport and analysis. As APOGEE progresses it will benefit on the focused work on database structures and access methods in the ALDAP project, and the PPDG studies of multi-site file caching in 2000-2001, that will run concurrently with it.

APOGEE intends to work together with the LHC experiments, and to contribute to the study and development of the major US Regional Centers for ATLAS and CMS[3]. The timing of APOGEE covers FY2000-2003, when these regional centers must be designed, prototyped, and their configuration optimized, in preparation for the installation, commissioning and initial use of these systems for "mock data challenges" between 2003 and 2005.

APOGEE is a crucial complement to the GriPhyN (Grid Physics Network) project submitted by the LHC experiments and LIGO to NSF, which proposes to implement 17 "Tier2" Centers, including 15 in the US and 2 at CERN, working in coordination with the major Tier1 Centers. GriPhyN focuses on the equipment and manpower to commission and operate the centers, and to use the tools of other projects, such as MONARC (1998-2000) and this project APOGEE (2000-2003), to provide the necessary tools for simulation, optimization, and strategy development. Should GriPhyN be funded, the manpower available for R&D studies would go to study the specific data analysis-related operations for the individual experiments, with a focus on the Tier2 centers.

---

[3] The exact scope of joint work will be determined through joint discussions among HENPDAG, US CMS, US ATLAS, the FNAL and BNL host laboratories, and the CMS and ATLAS computing groups at CERN.

The APOGEE instrumentation and simulation teams would carry the main responsibility for overall system optimization, covering the whole Grid hierarchy, and for maintaining a view towards optimizing the system in "standard" ways that could apply to future large experiments (e.g. not focusing too narrowly on particular platforms or software products).

Long Term Vision

We intend that the HENP Data Grid be one of the projects that will be the foundation of a longer-term plan, building on the PPDG and earlier projects, and leading into the production-scale Grid hierarchy that will include the Tier2 regional centers proposed in the GriPhyN project. In this long-term vision we identify three major phases for the Data Grid:

1. Support for distributed transparent read and write access to high volume data and computing resources
2. Support for a wide-area distributed computing system where data-intensive tasks move automatically to the node in the Grid at which they can be most effectively executed. We anticipate that "mobile agent technology" being developed offers a realistic and scalable way of matching the analysis tasks to the data;
3. Support for agent technology and "virtual data" to manage extremely large datasets. Virtual data are derived data defined only in terms of the input data and the process or agent that could create the derived data. Whether a virtual data object or file exists in many, one or zero copies is left to "Data Agencies" in the Data Grid to decide. The use of large numbers of mobile agents to replicate and manage data is inherently scalable to the multi-PetaByte datasets that will be generated in the LHC era, and by other future HENP experiments.

Substantial progress towards this vision will require major advances in network-aware middleware, in distributed and scalable metadata catalogs, and in the integration of scalable software configuration management into Grid services. Some of these developments are underway in the PPDG project now. The field-testing of the future services, in conjunction with the site-configurations and data analysis strategies to be developed in APOGEE, is planned to take place in the context of GriPhyN, once the first-round Tier2 centers have been installed. Stress testing of the system and further evolution of the APOGEE strategies will occur during the "mock data challenges" for LHC, and using real data from Babar, RunII, RHIC, JLAB and LIGO in a distributed environment.

The APOGEE project also will provide the forum for coordinating its own specific work plan with those of projects, such as the PPDG (which will provide essential software foundations) and the proposed GriPhyN project providing many of the physical components and the "Tier2" computing facilities of the Data Grid.

Manpower and Budget Analysis

The personnel required to allow APOGEE to meet its technical and coordination goals are detailed in the table below. We have assumed that a ramp-up occurs in FY 2000 and have set the FY 2000 requirement to exactly half that in FY2001. The estimates for FY 2002 and beyond must necessarily carry the caveats that:

1. APOGEE must provide any vital software components that will not be created by other Industry, Computer Science or HENP efforts;
2. APOGEE must take advantage of relevant Industry or Computer Science advances even though these may not be predictable more than months in advance.

	FTE FY00	FTE FY01	FTE FY02	FTE FY03	FTE FY04
Instrumentation					
Low-level data capture	0.5	1	0.75	0.75	0.75
Filtering and collecting agents	0.5	1	1	1	1
Data analysis and presentation	0.5	1	1	0.75	0.75
HENP workload profiling	0.5	1	0.5	0.5	0.5
Simulation					
Framework design and development	1	2	1.5	1	0.5
User workload simulation	0.5	1	0.75	0.75	0.5
Component simulations (network, mass-storage system, object database etc.)	1.25	2.5	2	1.5	1
Site simulation packages			1	1	1
Instrumentation/Simulation Testbed					
Instrumentation of existing experiment(s) (e.g.PPDG)	0.5	1	1	1	1
Acquire and simulate performance measurements	0.25	0.5	0.5	0.75	1
Acquire user workload profile	0.25	0.5	0.5	0.25	0.25
Test prediction and optimization			0.5	0.75	0.75
Evaluation and Optimization					
Quantify evolving needs of physics (including site policies etc.)	0.25	0.5	0.5	0.5	0.5
Develop metrics for usefulness of data management facilities	0.5	1	1	1	1
Optimize model systems			0.5	1	1.5
Long-Term Strategy (Towards "Virtual Data")					
Tracking and testing HENP/CS/Industry developments	1	2	1.5	1.5	1.5
Development projects in collaboration with HENP/CS/Industry where possible			0.5	1	1.5
Project Management (APOGEE and PPDG)					
Project leader (physicist)	0.5	1	1	1	1
Lead computer scientist	0.5	1	1	1	1
TOTALS	**8.5**	**17**	**17**	**17**	**17**

The personnel estimates are translated into dollars in the budget analysis below. HENP data management is proving an attractive area for Computer Science students to gain valuable experience and understanding. We expect to take advantage of this to keep the average person cost at $125k (including benefits and overhead) in most areas. In contrast, in the Project Management roles it will be essential to attract a physicist and a computer scientist with established and deserved reputations.

Appropriate use of commercial software components for instrumentation and modeling in addition to data management itself is an essential part of the APOGEE plan. This is the largest non-personnel component of the cost.

Limited funding for testbed hardware is included. The vast majority of testing will take advantage of the over $40M installed base of data handling systems within the US-HENP program. Most of the APOGEE-funded hardware will be tightly coupled to much larger HENP-funded systems. APOGEE will be able to negotiate short-term use of the larger systems in return for some access to APOGEE systems by HENP experiments.

	$k FY00	**$k FY01**	**$k FY02**	**$k FY03**	**$k FY04**
Manpower					
Instrumentation	250	500	406	375	375
Simulation	344	688	656	531	375
Instrumentation/Simulation Testbed	125	250	313	344	375
Evaluation and Optimization	94	188	250	313	375
Long-Term Strategy (Towards "Virtual Data")	125	250	250	313	375
Project Management (APOGEE and PPDG)	225	450	450	450	450
Commercial Software	100	250	375	500	500
Testbed hardware (in addition to parasitic use of production systems)	150	400	400	400	400
Workstations, M&S, Travel, Training	128	255	255	255	255
TOTALS	**1540**	**3230**	**3355**	**3480**	**3480**

Essential Role of a US-Based Grid for LHC and Other Distributed Computing in International Projects

Because of the higher network bandwidths expected to be available in the US compared to transoceanic links at any time, the proposed set of facilities will allow us to preview and develop effective strategies for the worldwide-distributed systems to be used from 2005 onwards. The existing US-CERN transatlantic network link, whose bandwidth is expected to rise between 2000 and 2005, will be used as an essential (bandwidth-limited) part of the Grid, where state-of-the-art “Quality of Service” (QoS) techniques will be required to manage the bandwidth. The Grid also will serve as production-prototype for the more capable systems that will be required to support the full-luminosity LHC physics program, starting in approximately 2008.

APOGEE, Data Grids and the Future

A Data Grid is a natural model for future scientific and engineering collaborations, future corporations, and future information-exchange extending to every sector of society. [Ref4]. As the ability to store, process and transport data increases exponentially with the advance of computing and network technology, agent-driven grid systems able to transparently and rapidly access, process and deliver data to sites throughout the globe will become feasible. An optimized, cost-effective grid could become a principal driving force behind the advance of the US and world economies, based on ubiquitous access to information with a richness that is unimaginable in today's World Wide Web.

Within the realm of scientific research, based on recent history, we can be assured that future HENP experiments and other large international collaborations will continue to stress the limits of information technology in all its forms. In the context of "Data Grids" this means that the need to use, refine and develop further generations of Grid simulation and optimization methods will outlive the need for any particular Grid implementation. The APOGEE proposal is the first aimed at fully addressing this class of problems, as the key element in a chain extending from the initial Grid ideas developed in Computer Science, to the first production-scale implementation of a worldwide Data Grid for high energy and nuclear physics.

REFERENCES

[1] R. Haupt and S. Haupt, 1998, *Practical Genetic Algorithms,* New York: John Wiley.

[2] W. Banzhaf et al., 1998, *Genetic Programming: An Introduction*, San Francisco: Morgan Kaufman; J. Koza, 1992-1998, *Genetic Programming I, II, III*, Cambridge: MIT Press.

[3] See Chapters 6 and 7 and the reference material in Reference 1.

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