

European Union — United States

joint workshop on

Large Scientific Databases



Annapolis, Maryland, USA, 1999 September 8-10

Supported in part by the
National Science Foundation

(Grant IIS-9910140)

and the

European Commission

(European Union Information Society Technology Programme)

Organized by

the Center for Advanced Computing Research at the California Institute of Technology,
and by the European Laboratory for Particle Physics (CERN)

this report compiled by the organizing committee:

Roy Williams

California Institute of Technology, USA

<roy@cacr.caltech.edu>

Paul Messina

California Institute of Technology and US Department of Energy

<messina@cacr.caltech.edu>

Fabrizio Gagliardi

CERN

<Fabrizio.Gagliardi@cern.ch>

John Darlington

Imperial College, UK

<jd@doc.ic.ac.uk>

Giovanni Aloisio

University of Lecce, Italy

<aloisio@sara.unile.it>

This document and supporting material may be found at

<http://www.cacr.caltech.edu/euus>

version 10/14/99

EU-US Workshop on Large Scientific Databases

<http://www.cacr.caltech.edu/euus>

Report from the EC-NSF Workshop*

I. Executive Summary

Science is continuing to generate ever larger amounts of valuable data, but we are in danger of being unable to extract fully the latent knowledge within the data because of insufficient technology. To address this we propose the establishment of an **Expedition Center**, a virtual “center”, hosted at multiple geographical sites, similar in scope and thrust to the US-NSF Science and Technology Centers. Features include:

- **unification** of information and knowledge management between US and Europe
- strong **leadership** and continuity of purpose
- **funding** in the millions per year
- **longevity** of 5 years or more
- **flexibility** to seize new opportunities quickly and to shift the agenda rapidly

The center would be a network of excellence in specific research domains, emphasizing trans-Atlantic teams, supporting both basic and applied research, with large-scale testbeds and large-scale demonstrations. It would have a strong education and outreach component. There might be four sites in the EU and US with independent funding for visitors, travel, and workshops. These regional centers could share the resources of independently-funded facilities to create large-scale demonstrations and prototypes; such sharing could be achieved by collaboration agreements or by rental. There would be liaison to other activities, for example the Framework 5 in the EU, the Grid Forums and Digital Library Initiatives in the US.

A crucial requirement for this kind of collaboration is **trans-Atlantic data communication** that provides high bandwidth, high availability, and low latency. We recommend a study to consider and cost the options in detail.

We recommend funding **application driven, multidisciplinary research**, with the creation of prototypes, testbeds and full-scale implementations. Such research should always be close to the needs of a particular community, preferably directly connected to scientifically-interesting research. Such a scientific community should be geographically distributed and international.

We encourage the creation and reconciliation of data object and **metadata standards**, but only in a strongly-defined, discipline-specific environment, and with enough funding to produce relevant and useful software, not just a report. Further work could define metadata semantics, discipline specific data dictionaries, information models for organizing metadata, and data models for describing data set structure. We encourage projects that use established, extensible metadata standards. Where an existing standard exists, new projects should use, subset, or extend one of these standards, or provide good reason for any decision to start afresh.

Interoperability projects should be encouraged, that begin with two or more existing scientific databases, preferably already catalogued and/or online, together with a good reason and mechanism for combining the data. We should then encourage the implementation of this federation.

We recommend specific research on the following aspects of **distributed and/or large databases**: data clustering and caching; data redundancy, dynamic summarization, and query formulation to allow machine optimization and brokering; splitting queries into separate, local queries and cost estimation of queries; parallel multi-dimensional access and search methods, approximate search methods, and data compression; load-balancing of computational work and data in distributed systems, replication of data among regional centers, protocols for high-speed, parallel dataflow, and protocols for real-time steering and control of running jobs.

We also recommend supporting work on **security** mechanisms, especially work that coexists with other security mechanisms. We recommend that creators of scientific databases should be encouraged to consider in advance the **preservation of the data**. Preservation description information should be associated with the digital objects being preserved. We recommend investigation of ways to standardize requirements for IT-courses in the **educational curricula** of the domain sciences with emphasis in data modeling and use of databases. We recommend exploration of the profound impact of databases and networking on the **process of science**, including publishing, peer-review, collaboration, and data ownership.

*This workshop was sponsored in part by the National Science Foundation, under the grant IIS-9910140 (PI: Roy Williams) awarded by the Information and Data Management Program of the Information and Intelligent Systems Division. All opinions, findings, conclusions and recommendations in any material resulting from this workshop are those of the participants, and do not necessarily reflect the views of the National Science Foundation.

2. Introduction

This joint workshop was set up under the auspices of the Joint European Commission / National Science Foundation Strategy Group that met in Budapest in September 1998. The meeting derived from a joint collaboration agreement between the EC and NSF, signed by Dr. George Metakides (Director of Information Technologies for the EC) and Prof. Juris Hartmanis (Director of Computer and Information Science and Engineering at the NSF). Some themes that were identified include:

- Digital libraries
- Human centered computing and virtual environments
- Large scientific databases
- Intelligent implants

This report expresses the conclusions and recommendations of the Workshop on Large Scientific Databases, held in Annapolis, Maryland, USA in September 1999. The purpose of the workshop was to develop a report to the funding agencies outlining a possible solicitation to the research community, with emphasis on joint European-US work on Large Scientific Databases. Before the workshop, each participant submitted a position paper (these are available at the web site <http://www.cacr.caltech.edu/euus>). The results of the position papers, presentations, and group discussion are summarized in this report. There were 12 participants from Europe and 12 from the US, and they are listed at the end of this report. The last section of this report describes possible funding mechanisms.

2.1 Scientific Databases

Through sensors, experiments, and computer simulation, scientific data is growing in volume and complexity at a staggering rate. The cost of producing the data is very high: satellites, particle accelerators, genome sequencing, and supercomputer centers represent information generation sources that collectively cost billions. The last decade has witnessed a thousand-fold increase in computer speed, significant increases in detector size and performance, a dramatic decrease in the cost of computing and data storage capabilities, and widespread access to high-speed networks. Without effective ways to retrieve, analyze and manipulate these data, that great expense will not yield the benefits to society that we might expect. In order to efficiently handle terabytes of data, one needs database engines with fast I/O speeds and advanced query engines, that can access geographically-distributed data. Most existing tools do not scale to terabyte or petabyte datasets—tools for discovery, subsetting, searching, filtering and analysis. This report identifies the research and development issues related to the creation and use of large scientific databases, and describes approaches for addressing the associated challenges, concentrating on those that can benefit from joint European-US work.

Scientific data is getting not only larger and larger — multi-petabyte archives are routinely discussed — but its complexity is increasing, meaning that the extraction of meaningful knowledge requires more and more computing resources. The benefits of e-mail and ubiquitous travel mean that scientific collaborations are getting larger and more geographically dispersed. These remote scientists, with a mixture of computing equipment on their desks, need catalogues and indexes of the data archive, the ability to select data objects, to define complex processing to be done, then choose how the results are to be returned to them. In addition, adequate authentication systems are needed to control access to the data across multiple security domains.

— NSF Workshop on Interfaces to Scientific Data Archives, 1998, Pasadena, California
<http://www.cacr.caltech.edu/isda>

Collection-Based Research

The typical way to proceed with analyzing large datasets, is to first generate some abstraction of the dataset in the form of features or summarization, for example the generation of monthly means over simulated climate data, or

automated cluster analysis to identify regions of high activity, or to find anomalies in the dataset. Scientists then use this information to guide them to regions of interest. In this exploration phase, they often wish to “drill down” in the dataset. This is done by specifying a subset of the dataset, either directly, or by the summarization features, repeatedly refining the focus. Once something interesting is found, there may be a long run of the computer for pattern-matching or other data-mining: to find other places in the dataset which “look like that”. Finally, the results of the search can be used to create new knowledge.

Just as simulation was added to the two traditional scientific cultures — theory and experiment — so collection-based research will be added. The experimentalists take data, catalogue it, calibrate it, and make it available to others on the Internet; the collection-based scientist will reduce, mine, and sift that data to make or break a hypothesis from the theorist. In this new paradigm, the raw or nearly-raw data is published in a way that would be impossible with only paper journals: not just tables and graphs, but rather ramified palaces of logic. These would be subsequently reduced and interpreted by other researchers: thus the person who takes the data may be different from the person who reduces it to small, palatable representations that can be printed on paper. In this way we allow more access by more scientists, and we also open more of the scientific process to examination by the scientific community. Publication can involve not just a table of results, but also a pointer to the online raw data and the exact specification of the chain of software that has been applied to it.

Already schoolchildren are happy with Internet-based research; as they mature into the next generation of scientists, they will expect a similarly-structured research environment. Collection-based and simulation-based research will become as accepted and important in science as the traditional cultures of experiment and theory. We must create that environment, and also put in place an educational curriculum that teaches young scientists how to use it.

Database Interoperability

Science has historically proceeded through fusion of experimental measurements, sometimes apparently conflicting, into a coherent unified theory. There is an urgent need to advance science not only by using, but also by fusing information from multiple sources, from multiple digital archives. The greatest leverage will come if the integration is not only within the lines of established disciplines, but across these lines. This idea of *interoperability* or *federation* of archives was an important theme at the workshop.

Application Focus

Sometimes well-meaning technology research projects fail not because the project was a bad idea or badly executed, but because the potential users did not use the results, perhaps because it did not fit their needs, perhaps because they did not know the research was being done. In the rest of this report, we advocate building software and other products that are of general use — not specific to only one scientific discipline — but to prevent development in a vacuum, we feel that there must be a strong application focus to ensure acceptance of the results.

2.2 Characteristics of this research

We recommend funding application driven, multidisciplinary research in collection-based data management, addressing a specific scientific problem. **We recommend** the creation of prototypes, testbeds for real users to do science, that lead to full-scale implementations. Once the testbeds and infrastructure are in place, other applications and technology developers may be solicited. The work will consist of testing individual pieces, integration of pieces, prototyping; all this with applications that people care about (not toy applications). Distributed computing will be an important component, working with data, computing and clients spread across the Atlantic.

Domain scientists typically consider their work done once they have developed (or adopted) models to generate data. As a rule, they do not have the discipline-wide support needed to organize data using mutually agreed terms. Consequently they are happy to just work with files, not databases. Any database that will be developed for them

after the data start to pour in will not be used, unless the discipline has standardized information models and metadata attributes. **We recommend** investigation of how to get domain scientists to use databases as a rule rather than as an exception, perhaps looking at successful projects and the profiles of those who routinely apply databases. User-oriented research projects, with tightly coupled interdisciplinarity as a prerequisite, should be given a high priority.

The impact of large scientific databases is not just technical challenges such as enabling faster analysis. There will be a profound impact on the way that the process of science occurs, such as publishing, peer-review, collaboration, and data ownership. **We recommend** that these issues be explored also within the Expedition Center (Section 9).

Interoperability of scientific databases will drive global collaboration in the scientific and, later, in the commercial worlds. **We recommend** funding projects that begin with at least two existing scientific databases, preferably already catalogued and/or online, together with a good reason and mechanism for combining the data. We encourage the implementation of this federation.

If we expose real, non-public, scientific data to the open Internet, then there must be sufficient access control to assure the safety of proprietary data. **We recommend** supporting work aimed at establishing policies and mechanisms that provide this access control, especially work that coexists with other security mechanisms. **We also encourage** mechanisms that control the transition of proprietary data into public use.

Whenever scientific data is made available to a larger group, there must be supporting material to show its parentage and meaning. There must be sufficient documentation to show how the data was gathered and processed, so that (in principle) the data could be reproduced.

2.3 Benefits of supporting this research

Scientific data is already going online and public at an increasing rate, and the general public (who pay for much of it) should gain the benefit. There are web sites of genome information, so an educated consumer might combine her genetic profile with a Web surf through gene libraries to determine predisposition toward adverse drug reactions, for example, or for Alzheimer's disease, colon cancer or other afflictions that might eventually be treatable through gene therapy. As another example, a consumer might use a geographic system to evaluate a new area before relocating there: she might make a customized map of the area, with overlays of climate, crime statistics, earthquake hazard, and school coverage.

Many aspects of our lives, including food and airplane tickets, are cheaper because of the databases and networks that streamline delivery and reduce waste: in the same way, complex machinery, for example automobiles and electronics, is cheap because the computers allow multidisciplinary teams to work effectively. In the new global economy, high-performance database and workflow technologies will become even more critical. We believe that research in scientific databases will be twofold: as a foundation for industrial applications, but also as a seminary for the graduate students and postdocs who will innovate in this new economy.

One obvious benefit of EU-US cooperation is to make scientists more productive, since hardware design and software only needs to be done once. When a good design has been found and learned, it can then be reused on other projects. Once a scientist has learned an interface with one dataset, she is familiar with that interface when it is used in another place. In fact, much of the technology already exists, but needs to be integrated and used; it will save money by fully deploying what has already been developed.

Interoperability of databases is a major theme in this report. This means the fusion of diverse data to create new knowledge — for every pair of data collections that can interoperate, there is the chance of a consequent scientific discovery — thus interoperation of digital libraries creates knowledge and value without the expense of gathering more data. Interoperability means not only the ability of a group of researchers to share knowledge, but also the

concept of universal access to data, by researchers, educators, students, and the public. All of these communities should have access to current data that drives scientific discovery.

The Expedition Center should seek endorsements from industry, for example database and networking companies. We expect them in turn to benefit greatly from this research and development. However, we feel that such industrial endorsements should not be a prerequisite for funding, but rather that domain scientists should be the primary target partners.

2.4 How is this Different

Digital Libraries

Another part of the EU/US Joint Collaboration agreement concerns Digital Library (DL) technology. Here we consider its relationship to the topic of Large Scientific Databases (LSDB). Both are interested in the impact of the web culture and infrastructure, both are interested in strong, non-invasive security. But in many ways they are complementary. Digital Library technology has focused on information discovery, while Large Scientific Databases have focused on the ability to access and analyze large data sets.

- DL emphasis is on archiving of human-accessible data such as text, audio, and video, where the challenges come from indexing semi-structured or free text. LSDB are more concerned with abstract data that cannot be readily understood without choosing transformation, data-mining and visualization tools.
- DL are more interested in being accessible to those who are not computer professionals, whereas LSDB strive to provide fast and powerful data handling services to those who know exactly what they are doing.
- DL are interested in interpretation, publication, multi-lingual issues, and parentage; but in the LSDB, we want fast networking, caching and clustering technology, and query estimation.
- In the DL community, archiving and preservation includes scanning paper and extracting information from catalogues that were created as cards; for the LSDB, however, we consider how often tapes have been read, how to pack more data into a tape, and how to create a lasting information model for scientific endeavors, even for those that have not yet begun.

— NSF Digital Libraries Initiative Workshop Series
<http://dli.grainger.uiuc.edu/natlsynchpubs.htm>

Human-Centered Computing and Virtual Environments

Another workshop in this series of joint EU/US collaborations concerned Human-Centered Computing and Virtual Environments (HCCVE), which considers how people can apprehend large or abstract datasets, how they can work effectively with others via computer, how to design new hardware that can be used effectively by many communities. This research in LSDB is the perfect complement to this theme, which concentrates more on technical and scientific issues, on how to create, discover, store, and gain insight from the data that is presented by the HCCVE.

Commercial and Industrial Databases

Obviously we do not want the Commission and the NSF to subsidize the enormous industry that surrounds commercial databases. Thus **we do not recommend funding** servers for standard web pages, or databases tuned for e-commerce transactions. However, scientific databases are different from current commercial offerings in enough ways, as discussed below, that we feel government funding to be appropriate. Indeed, we hope that the technologies identified and seeded by this report will be not only of scientific value, but commercially profitable to European and US businesses in the years to come.

Commercial databases and data warehouses tend to concentrate on the rapid analysis of current information, with all

data stored on rotating disk. For the largest corporations, the data warehouse can be over a hundred terabytes in size, comprising billions of records. Scientific data collections can be even larger, with petabyte sized archives being planned. Furthermore, scientific data objects are often quite large, large enough to require compact storage schemes and streaming interfaces. Scientific databases often have a dichotomy between online and nearline storage. Online data can be accessed in an interactive way, whereas nearline data may be in a robot-controlled tape archive that takes many minutes to get to a given data object; access mechanisms involve batch jobs, data staging, job status inquiry, or interactive steering.

Data-handling systems for scientific data must be very flexible, since it is often not known how the data will be used when the system is being designed. Scientific users often want to write their own programs and run these on a server; besides security problems, designers must try to accommodate a variety of languages and a variety of levels of expertise.

The scientific community, with its focus on research, has pioneered many concepts that were eventually accepted by the commercial world, and we expect the same to be true with the complex database and networking technology envisioned in this report.

2.5 Examples of Large Scientific Databases

This report discusses and recommends initiatives in Information Technology that might be undertaken jointly by the EU and the US; we are not recommending discipline-specific initiatives. However, it seems appropriate to provide context by describing some scientific disciplines and the way in which they use, or could use, large databases.

Furthermore, to make sure that the research stemming from this report is relevant, **we recommend** that it always be close to the needs of a particular community, perhaps directly connected to scientifically-interesting research. To be a candidate for a prototype or testbed, such a scientific community should preferably be geographically distributed and international.

- Scientific Data Archives at Caltech
<http://www.cacr.caltech.edu/SDA>
- Scientific Data Analysis at the University of Maryland
<http://www.cs.umd.edu/projects/hpsl/ResearchAreas/research2.htm>

High Energy Physics

A data thunderstorm is gathering on the horizon with the next generation of particle physics experiments. The prime data from the next-generation CERN CMS detector will amount to over a petabyte (10^{15} bytes) per year that is to be archived, with subsequent analysis to find rare events resulting from the decays of massive new particles. Within the next decade, CMS and other experiments now being built to run at CERN's Large Hadron Collider expect to accumulate on the order of 100 petabytes. Object-oriented database systems form the foundation of choice for many of these systems, and there is much interest in the creation of globally-distributed regional repositories that appear to users as a unified whole. Because each experiment is centrally managed, the political problems of information modeling for interoperability are less severe than for other scientific communities; however, there is always a risk that different experiments will choose different and incompatible systems.

We encourage projects that support interoperability with other scientific disciplines to promote the use of the advanced technology, but perhaps narrow focus, of the high-energy physics community.

- The Particle Physics Data Grid
<http://www.cacr.caltech.edu/ppdg>
- MONARC: Models of Networked Analysis at Regional Centres (CERN)
<http://www.cern.ch/MONARC/>

Astronomy

Astronomy will experience a major paradigm shift in the next few years, driven by large, systematic sky surveys at multiple wavelengths. We believe that these digital archives will soon be the astronomical community's main avenue for accessing data. Systematic exploration and discovery in these databases will play a central role in the day-to-day research activities of most astronomers.

An example of an emerging international collaboration is the astronomical Virtual Observatory (VO). In the past, astronomy proceeded through observation of individual celestial objects, in the same way as a medical doctor works with individual patients; however a new paradigm is emerging because of the existence of digital sky catalogs: scientific databases that will be tens of terabytes in size. Astronomers can now work statistically, with *collections* of celestial objects rather than individuals, rather as an epidemiologist works with populations of people. Just as epidemiology provided a new dimension to the study of health and disease, so we expect the Virtual Observatory to result in new astronomical discoveries.

Astronomers will be able to launch automated multi-wavelength search and discovery among all known catalogued astronomical objects. There will be powerful and novel data analysis and exploration environments for a qualitatively new and different type of astronomical research: multi-wavelength exploration and discovery over the entire sky using all known catalogued astronomical objects simultaneously. The discovery process will be accelerated through the application of advanced visualization, data mining, and statistical tools.

Astronomical data has the advantage of relative openness, so that security is less significant. However, the different wavelength regimes (radio, optical, X-ray etc.) have traditional cultural differences amongst the researchers that must be respected in the creation of information models.

- National Virtual Observatory
<http://www.srl.caltech.edu/nvo>
- AstronomyData and Archive Centers
<http://cdsweb.u-strasbg.fr/astroweb/center.html>
- Astronomy Digital Image Library at NCSA
<http://imagelib.ncsa.uiuc.edu/imagelib.html>

Gravitational Wave Observatories

An existing EU-US collaboration links the US LIGO and the French-Italian VIRGO gravitational wave observatories; these are large government-funded facilities designed to detect astrophysical sources of general-relativistic gravity waves (GW), primarily from the coalescence of massive objects such as neutron stars or black holes. The detectors produce hundreds of terabytes of data per year, and the extraction of astrophysically significant signals will require not only high-performance computing, but also integration of other, external, data streams. The discrimination of local events from astrophysical signals will obviously require continuing strong collaboration between EU and US observatories. Looking to the future, if GW sources are to be localized in the sky, all observatories (there is also a Japanese project) must be even more tightly coupled.

- Laser Interferometric Gravitational-Wave Observatory (LIGO)
<http://www.ligo.caltech.edu/>
- VIRGO Gravitational Wave Observatory
<http://www.virgo.infn.it/>

Remote-sensing and Geographical Systems

A highly significant range of scientific data is geospatial data — that which is associated with a region on the surface of the Earth. There are traditionally two communities: vector-based and image-based. Vector-based mapping is a continuation of the hand-drawn maps of the past, made modern by computer technology and now called GIS

(geographic information systems). Image-based mapping comes from the space-science community, however, from remote-sensing satellites. Effective combination of these cultures is moving quickly, but we recommend further projects that might join these cultures. In these projects the issues related to the distributed geolibraries must be solved in order to integrate the functions of a friendly web browsing with those of GIS and related technologies. **We encourage** projects that integrate maps and images of the Earth's surface. Distributed data and computing facilities must be designed for data integration, for secure access to the distributed information, and for flexible management and control. In this way, the full benefits of space-based imagery will be available to a larger segment of society.

Large quantities of Earth-observation data have been acquired in many wavelengths by remote sensing space missions, and accompanying software systems have been built by the National Space Agencies. Techniques such as interferometric SAR and hyperspectral imaging provide increasingly comprehensive views of our planet. **We recommend** the funding of projects that provide interoperability and unification of the ESA and NASA collections; also interoperability between remote-sensing and other fields, such as archaeology, ecology, seismology, and geology; also new ways to use and combine existing data objects in new ways.

The designing of distributed geolibraries requires the agreement on suitable metadata standards. The geospatial community already has some metadata standards that have evolved over years, such as the FGDC and CEOS standards. We recommend that new projects in the geospatial domain should use, subset, or extend one of these standards, or provide good reason for using something else.

- Workshop on Distributed Geolibraries (NRC) held in Washington on June 15-1998
<http://www.nap.edu/catalog/9460.html>
- NASA Global Change Master Directory
<http://gcmd.gsfc.nasa.gov>
- Alexandria Digital Library
<http://www.alexandria.ucsb.edu/adl.html>
- NASA Earth Observing System and Data Information System (EOSDIS)
http://spsosun.gsfc.nasa.gov/New_EOSDIS.html
- Committee on Earth Observation Satellites (CEOS)
<http://gcmd.gsfc.nasa.gov/ceosidn/>
- Intelligent Satellite Data Information System at DLR
Deutsches Zentrum für Luft- und Raumfahrt
- National Environmental Satellite, Data, and Information Service at NOAA
<http://www.nesdis.noaa.gov/>

Bioinformatics

In the last ten years, the new field of *bioinformatics* has formed a bridge between biology and information technology, including computational analysis of gene sequences. New experimental techniques, primarily DNA sequencing, have led to an exponential growth of data, leading to a change in emphasis from accumulation of data to its interpretation. However, the computational tools for sequence analysis, while evolving rapidly, are still not sufficiently scalable for these large sizes. There is a cultural gap between the traditionally “wetware” orientation of the biological community and the computer scientists; to make matters more difficult, much of the bioinformatics research occurs behind the closed doors of commercial interests.

The data can be broadly categorized into sequence, expression, and protein data. A sequence represents a DNA or RNA molecule: the “program” from which proteins, and hence organisms, are created. Protein data represents the biochemical and biophysical structure of the molecules of which the organism is made. Expression is the least-known part of the machinery: how different parts of a sequence are activated (“expressed”), causing the creation of proteins.

Candidate initiatives in bioinformatics are:

- Creation of interoperability mechanisms by which bioinformatic data can be exchanged and fused; for interoperation of existing databases, for interoperation and fusion of sequence, expression, and protein databases; for interoperability of data that relates to different biological species.
 - The creation of a “genome clearing house” — a digital library — for public-domain biological data. This central repository, managed jointly by EU and US authorities, would act as a *de facto* creator of standards, as well as giving maximal leverage to the genomic data that is public. This model has been successful in the discipline of nuclear physics: the US National Nuclear Data Center, and the data services of the OECD-funded Nuclear Energy Agency are good examples.
- US National Nuclear Data Center
<http://www.nndc.bnl.gov/>
 - OECD Nuclear Energy Agency
<http://www.nea.fr/>

Simulation Output

Computer simulation has gained great stature in recent years, as an adjunct, sometimes even a replacement, for experiment and theory. In aerospace design, for example, hypersonic flow is a very difficult regime for wind tunnel experiments, so simulation is very important. In most astrophysical regimes, such as galaxy formation, it is obviously impossible to do experiments!

In the past, it has been a single person or group that created and ran the scientific simulations, then reduced the results to a peer-reviewed paper for publication. As collaborations get bigger and more dispersed, we see the need for the simulation output itself to be available to collaboration members as a digital library, complete with catalogs, finding aids, visualization tools, as well as the big data itself. The same structure is needed when comparing simulated with experimental results, except that an additional requirement is one of interoperability, so that experimental and simulated results can be compared with minimal effort and with unified, published calibration.

Efforts in this direction are underway in Europe and the US, and **we recommend** the encouragement of projects that can turn simulation results into a digital library in this way; especially interesting are those projects that promise interoperability with external systems.

Industry will benefit greatly in future years. Large projects such as design of a new aircraft will have multinational teams and a distributed data infrastructure for engineering and workflow; it will be essential to have simulation tightly integrated with the framework.

2.6 Roadmap

This report details the state of large scientific databases, and what joint EU-US research could be initiated to make database technology more useful to scientists and to the general public. The rest of the report is divided into the following sections:

- Scalability: as the price of storage and networking falls, the size and complexity of databases increases; but other aspects of the database must keep up if it is to remain useful to the end-user.
- Information modeling: What is the nature of the contents of the database and its catalog? Users must understand this before they can use it; also it forms the basis for interoperability between data collections.
- Interoperability: How can information from multiple collections be fused to extract new knowledge?
- Information Flow: How does information move in a complex system, how do users discover the database and its capabilities, how do the users initiate and control a complex processing pipeline?
- Preservation of databases: How to ensure that digital scientific data is still available, when necessary, many years in the future.
- Education and outreach: This chapter addresses ways to educate the next generation of scientists in the

methods of collection-based research.

3. Database Scalability

It is not very difficult to put some scientific data into a file system, and use some basic office software to work with it. For datasets of a few megabytes, there are many solutions available. Besides, upgrading is simple: binary is converted to text, and all the data can be ingested into a new scheme. Summaries can be created on the fly and the catalog fits into a single web page. But when the database grows, things get more difficult:

- Scalability with respect to the quantity of bulk data in the database. Numbers in the multi-petabyte range are considered in some scientific disciplines (high-energy physics), and these sizes are also being considered for databases of simulation results from supercomputers.
- Scalability with respect to geographical separation of the database components: in other words, networking and communication requirements. Here we also include scalability with the number of databases that are to be joined into a single, unified virtual database.
- Scalability with respect to the size of the user community
- Scalability with respect to the defining limits of applicability of the database: the organization of the database needs to be extensible. For example the frequency range of a physics experiment that records electromagnetic radiation may be extended when new instruments or detectors are added to the experimental apparatus. The data collection will have to track this change in attributes and data structure.
- Scalability with respect to the duration of the database project, or in other words the behavior of the database as technology evolves, as user requirements changes, as funding structure changes, as time advances.
- Scalability with respect to the complexity and heterogeneity of databases to be federated.

Some of these aspects of scalability will be discussed below.

3.1 Infrastructure for Large Scientific Databases

An architecture for a large scientific database may be as follows. A set of regional centers hold subsets of the full data, with the subsets overlapping for the most popular data. The regional centers may have idiosyncratic implementations and may be independently curated, but there is a uniform set of protocols and services available at each center, so that each “looks the same” from outside. The centers are connected by high-speed, perhaps dedicated networking. Users see a single entity: they are not aware of the distributed nature of the database. The data collections are backed up by archives that support disaster recovery. Data movement is managed by a data handling system that understands replication of data between the collections. Information discovery systems provide flexible access to the data, logically decoupling the storage of the data from the access mechanism.

3.2 Database Size

A petabyte is 10^{15} bytes: enough to store a few high-resolution images for every human being on Earth, enough to fill a rail-car with high-capacity tapes. Streaming this quantity of data through a single high-speed tape drive would take over 3 years, so the creation of such a dataset must be done carefully, since it is so difficult to see every byte, for example for recalibration or computing summary data.

In the next few years, working with these quantities of data will be a challenge to industry as well as science, and the solutions from this research will be more widely used. **We should encourage research on:**

- Hierarchical storage systems, so that more frequently-used data can be kept on faster media, and backups and replication can be automated. Performance is improved by segregating large and small files. Data services can be easily implemented through a comprehensive API.

- Distributed storage systems, integrated by a single virtual data-handling system that connects geographically-separated components. The storage systems need to be fast, reliable, secure, and replicated, while allowing multiple administrators.
- Parallel data delivery, so that a program asking for many data objects can be spawned on multiple processors, and receive the objects in parallel and asynchronously.
- Interoperability of “big data” systems, so for example tapes can be written with one system and read by another.

Requests that use data that is spread around the system are much more difficult to service than when the data is all on the same physical storage unit. For example when a large table is stored row-by-row, it is easy to extract a single row, since all the elements are together, but it is much less efficient to get a column, since sparse access is needed.

We recommend research areas on these issues:

- Clustering, the question is which data objects should be stored “near” similar objects,
- Caching: the question is which data objects should be on fast storage,
- Redundancy: the question is which datasets should be stored redundantly in different organizational patterns,
- Indexing: creating efficient ways to search scientific data (that are not addressed by commercial systems),
- Summarization: the question is when should summary data be computed on-demand, and when pre-computed.

3.3 Networking

A crucial requirement for effective EU-US collaboration in large scientific databases is trans-Atlantic data communication that provides high bandwidth, high availability, and low latency. The most important metric is **throughput**: how much data can be moved across the ocean in, say, 24 hours. This may be a small fraction of the nominal bandwidth. Without this link, many significant projects simply cannot begin. The link may be dedicated and government-funded; it may be a commercial venture; or it may be some combination. **We recommend** a study to consider and cost these options in detail.

Regional data centers communicate with each other differently from the way they communicate with users. An individual user, making an *ad hoc* request to the system, does not require high bandwidth, but does require the low latency for good interactivity. On the contrary, regional centers interoperate through bulk transfers that are predictable in advance, transfers that require high bandwidth, but can tolerate high latency. These transfers provide support for updates, backup, or replication, and to make new data products through joins or other cross-references. A hybrid situation is when a user makes a large request of one regional center, but the data is actually at another center. In the future, *quality of service* will be important — the ability of a channel to deliver data at a uniform, guaranteed rate, for video and audio applications, or for replication of a real-time data source. Another aspect is the possibility of scheduling bandwidth for a computer run or for video/visualization conferencing.

Streaming Data

Since the sheer amount of data is a dominant factor, we are led to a new paradigm for analysis: scheduled streaming. When a program needs to access a very large amount of data, the traditional approach is that the program requests data objects in an arbitrary, probably suboptimal sequence. In place of this, we might think of putting the program in a queue so that next time the collection of data is streamed through the system, then that program, as well as many other programs, will see each data object. In this way control is handed to the data system, which is much better able to optimize the process than the end-user scientist. These ideas are being actively explored in a number of existing projects.

To move in this direction, we could consider data architectures where the emphasis is data movement rather than data storage. Such a shift raises numerous issues about how data management systems should be constructed, including:

- Shifting from file-oriented to stream-oriented processing. Processing models that assume a data source has an “end” aren’t suitable for data streams that continue indefinitely. Incremental processing will be more the norm, and splitting and combining of streams will be important operations.
- Constructing new kinds of data management components. While most of the basic system components in current database management systems will continue to be of use, there are other functions that don’t appear in current systems. One such component is an alerter, which processes a data stream against thousands of stored conditions. Another is an accumulator, which is a high-turnover storage manager for maintaining a shifting subset of data that has appeared in a stream.
- Alternative structures for data. Normalized, tabular forms of data may not always be a suitable representation, as they can place related information far apart in a data stream. Hierarchical, grouped structures, as typified by XML, may be more appropriate.
- New roles for metadata: it will need to be mixed intimately with the data it describes and annotates, for purposes of routing and assessing relevance of data in streams.

3.4 Distributed Databases

When diverse, *ad hoc* queries appear in a globally-distributed database, it may be that data is to be moved from tape to disk, from machine to machine, from place to place, and across the ocean. The success of a large distributed scientific database depends crucially on the optimization of this process.

We recommend research work on how queries and processing requests can be formulated to streamline this optimization process; also on how such a query can be split in separate, locally-executed queries, with machine-specific data access. **We recommend** work on how the cost — in terms of computation, communication, and time — can be estimated before and during execution. Fruitful work could also be done by evaluating projected and actual use-cases of distributed database projects.

We recommend research into ways of specifying, estimating, and executing complex access patterns, especially when the data is widely distributed, especially when the meaning of the data is from different scientific cultures.

We also recommend research into

- Load-balancing: how computational work and data are spread around the distributed systems,
- Replication: what should be replicated among the regional centers,
- Protocols for high-speed, parallel I/O, with investigations of synchronous and asynchronous delivery,
- Protocols for real-time steering and control of running jobs.

3.5 Scalability in Complexity

Example: Gravitational Wave Observatories

Currently, databases from different scientific cultures are joined in twos and threes: however, in the future we might expect many such databases to interoperate. As an example, consider the LIGO and VIRGO gravitational-wave observatories, a collection of large scientific instruments in the USA and Europe for detecting astrophysical events that involve compact objects such as black holes and neutron stars, run by different funding agencies and separated by thousands of kilometers. These instruments are plagued by noise sources: seismic, acoustic, magnetic, cosmic-ray, gamma-ray bursts, and so on; and no astrophysically-significant signal can be found without careful discrimination and veto of noise. To do this fusion effectively, we must have access to these diverse data sources and merge them in a usable form for the data-analysis software.

4. Information Modeling

In the past, a major question with scientific data concerned *file formats*, and in this new age of global internetworking, that question is transformed and extended. There are new ways to encapsulate such files and to extend them with the addition of metadata, using MIME and XML technologies. At the same time, distributed object systems such as CORBA, Java RMI, and Voyager are allowing machines to exchange objects directly between trusted systems. Authentication and authorization have become serious issues. **We should encourage** projects which explore these software advances, so long as they have a context in application-driven projects.

An information model is the generalization of a file format with the following implications:

- Data objects appear in hierarchical collections, so that a complex object can be made up of a collection of simple objects, which can in turn be members of other collections.
- There is metadata associated with the data object, as well as a simplification of the data content to allow browsing and summarization.
- An object can inherit from a simpler object (all *galaxies* are also *celestial objects*). This allows the creation of sharp, specific tools for a particular purpose, while reusing all the software from the parent class.
- An object can be defined to be a member of a generic class, thus allowing generic tools to act on specific types (a *galaxy* may be *drawable*). Thus we can bring existing software to bear on new types of data objects.
- There may be a validation procedure, to check if the object is syntactically and semantically valid (if this object is supposed to be a *geographical latitude*, is it between -90 and 90 degrees?).
- Data objects may be static (a file) or dynamically generated from a query.

In a distributed system, emphasis shifts from objects to services. A request object is sent to the service, and a response object is returned. Potential users of the service need to know that it exists, presumably through a discovery service, they need to know how to use it — how to construct a request and what kinds of response objects are available. Services should be designed for use by either a human or a machine, meaning that the response may be cast as a structured document that the (machine) client can interpret.

— XML: Extensible Markup Language
<http://www.w3c.org/xml> and <http://www.xml.com>

Example: Federated Astronomical Sky Surveys

Already there are sky surveys, databases of celestial sources and images, each maintained and curated by a dedicated group of astronomers. Each group has independently decided on the nature of objects, services, and interfaces. To make the virtual observatory described above, we need to make standards for these things that are flexible, open, and unrestrictive, yet powerful and comprehensive.

4.1 Standard Scientific Data Objects

“There’s no problem with standards: as long as you use mine!”

We want sufficient standardization of the scientific data objects so that the individually curated parts of the distributed database can work together effectively. On the other hand, we do not want to stifle creativity and flexibility with rigid bureaucracy and overambitious standards. The characterization of information in this way is itself a research topic, requiring the development of a common information model. The broader the community that develops the information model, the sooner we will be able to support information exchange within disciplines and between disciplines. **We recommend** funding for the creation and reconciliation of these standards, but only in a strongly-defined, discipline-specific environment, and with enough funding to produce relevant and useful software, not just a report.

Support for information models is rapidly evolving. Within the last two years, the use of semi-structured representations for information has made the characterization and sharing of data much easier. Many scientific collaborations are using the eXtended Markup Language (XML) and its schema definitions to tag the structure and semantic context of their data sets and data collections. **We encourage** projects that use such naming mechanisms, especially when they use an established, extensible metadata standards such as the Dublin Core or RDF (Resource Description Framework). **Joint collaborations** will facilitate extension of these initial efforts to develop interoperable metadata semantics, discipline specific data dictionaries, information models for organizing metadata, and data models for describing data set structure. Access tools based upon the information and data models will then be needed to facilitate use by the public and educators.

An emerging need within science communities is an information model that is independent of the operating system and storage system. An infrastructure independent representation will facilitate federation of data collections, and migration of data collections forward in time. Such a representation needs to build upon the current data models that are being developed in the joint EU/US collaborations, and should also include collaboration with the computer science community. The goal is to create information and data models that can then be used by other scientific disciplines.

The impact of such collaborations can be immense. The integration of discipline specific models can lead to standards for metadata attributes, and for metadata representation. Given the development of standard representations, it is then possible to integrate information across multiple disciplines, and improve the level of science that is conducted. The ability to exchange information between disciplines is the fundamental basis for interoperability between disciplines, and is thus the core requirement for inter-disciplinary research. By identifying the needs of multi-disciplinary research, the common infrastructure needed to support science in general can be understood. Thus the effort devoted to joint applications can serve as the driver for development of common infrastructure.

- Dublin Core Metadata Initiative
<http://www.oclc.org/oclc/research/projects/core/index.htm>
- Resource Description Framework (RDF)
<http://www.w3.org/RDF/>

Publishing Information Models

It is possible to publish the schema used to organize a collection as an XML schema. Information discovery can then be achieved through queries based upon the semi-structured representation of the collection attributes provided by the schema. Distributed queries across multiple collections can be accomplished by mapping between the multiple schemata. The context to associate with a data set can be specified by defining the data set structure, its physical interpretation, and its associated origination information. Since the organization of the attributes can be defined through a schema, it is possible to archive the information about the collection independently of the data sets that comprise the collection. This makes it possible to migrate a collection forward in time onto new technology. The collection description is instantiated on the new technology, while the data sets remain on the physical storage resource. Or conversely, the data sets are moved to a new physical storage system, while the collection description remains on the original system.

Object Oriented Approaches

The use of object-oriented languages and object persistency is becoming ubiquitous in scientific data processing: these technologies allow us to define, implement and store the complex science objects and inter-relationships that we deal with. We can then express highly complicated queries on the object store in order to extract the events and features of interest. **We recommend** the exploration of information models that have object-oriented characteristics of extensibility, so that the model is a serialization of the object itself.

A promising system for robust, flexible storage of large (100 terabyte) data is HPSS, which is a file-based system. But there is a mismatch when it is used to store an object-oriented database: one interesting issue to pursue is the development of data models integrated with mass storage models to preserve performance (in term of access speed) and scaleable sizes.

- High Performance Storage System (HPSS)
<http://www.sdsc.edu/HPSS>

5. Database Interoperability

Essentially every scientific discipline is developing large data collections and the mechanisms to support access to the collections. Now is the time to unify the efforts through the creation of common infrastructure, thereby providing interoperability between European and US scientific databases. This implies common interfaces — discovery, data handling, remote procedures and storage resources, and also a common information model — applicable to data, collections, and procedures, as well as semantic interoperability.

Properly designed collaborations provide not only the ability to integrate scientific activities within a discipline, but also provide access to technology that is being developed within other communities; development of information management technologies is an excellent target for better integration of computer science and computational science communities.

Interoperability is an effort that **requires explicit funding**, otherwise local solutions are developed. Local solutions imply restrictions on support for a single discipline, or a limited set of data collections. There is an evolutionary path that can be followed to build information resources that are available to the larger community through the federation of collections. The current approach is to install wrappers in front of existing collections that transform the information content into a standard representation. This enables the distribution of queries across legacy systems for the retrieval of information. A similar approach is used for supporting data access to legacy storage resources. Wrappers or servers are installed in front of the storage systems that support access through a common API. While wrappers provide an interoperability capability, they tend to be limited to the manipulation of relatively small data sets. Large scale data manipulation requires the tight integration of data and compute resources.

Example: Satellite remote sensing

Suppose there is a study of environmental damage, of how a delicate natural habitat is changing. We have satellite images from one and two years ago, with which we can apply a classification algorithm to establish the area of the habitat and how it changes with time. By using finding aids, we discover that there are satellite images from five years ago from the space agency across the Atlantic. We would like to run the same classifier, and extend the study in time, thereby making the result, and its political impact, much stronger.

Currently this kind of interoperation is very difficult: ESA and NASA have different catalogs, different storage mechanisms, different authentication and access policies, and it is difficult to interoperate in the manner described above. However, the value of these expensive data collections can be greatly enhanced by agreements and discussions between the data curators as detailed below — a high leverage from small funding.

5.1 Open, Federated Architectures

Based on our recent experience, **we would encourage** open architectures for federated scientific information systems in the following senses:

- Resource sharing with minimal administration cost — organizations can make their resources (data and programs) available to the federation without going through complex installation procedures, independently

of the local hardware and operating system.

- Operational autonomy and security — participation in the federation does not affect (in an uncontrollable way) the ability of the organization to provide services to its local clients; other members of the federation cannot access resources that have not been explicitly exported.
- Efficient resource discovery and access — users can locate resources quickly and accurately and they can access them in a straightforward way, via a network-transparent interface, almost as if they were available locally.

The importance of making the system open and flexible while keeping its administration cost low must be stressed as it significantly affects its acceptance. This is especially true for research environments where software development and installation is primarily done by scientists, not by a dedicated task force as is the case in a business environment.

While these criteria assess the extent to which data collections will take part in a federation, we must also consider the likelihood that the target group of end-users will use it. **We should encourage projects** with the following ingredients:

- Multiple, independent databases, preferably curated by different groups of people, preferably in different countries, preferably culturally separated,
- Scientific justification for federating the databases,
- An architecture that encourages access by a set of people that is more than the architects of the system.

5.2 Metadata for Interoperability

Metadata means “data about data”. Catalogs, indexes, directory structures are all metadata. The role of metadata is to bring diverse resources together by capitalizing on their commonalities: for example, librarians work with books, each of which has the same basic *schema*: each has a title and at least one author, a publisher, date, etc. In a similar way, we would like to identify common scientific data objects, perhaps a time series, a region of the sky, a gene sequence.

- The broader the community the metadata is meant to serve, the less detailed it can be,
- How one repository applies a schema to its resource may be somewhat different from another repository,
- Many details of a repository may not be visible from an external view.

Thus, it is important that a first-order framework allow the user to “drill-down” to the repository-specific information and services.

Bringing together a broad community can lead to large metadata schemas, and large schemas are difficult to maintain and intimidating to implement, for example with Z39.50 search profiles, the BIB-I metadata schema has over a hundred concepts, and GEO-I over three hundred.

One approach that can address these issues is to break the metadata standard into small pieces. Dividing the standards by discipline or community is perhaps the most natural way to differentiate metadata; however, certain classes of metadata (e.g. bibliographic) may span many communities. Projects are underway to consider a community as a hierarchy of sub-communities; it is natural then that the division of metadata standards should be hierarchical in order to represent varying levels of detail. For example, a general profile could be used for searching across all of space science, while sub-profiles for the component communities (planetary science, space physics, satellite tracking) would be defined for searches within a narrower disciplinary range.

This hierarchical, community-based approach has several advantages. First, each sub-schema can be kept small. Responsibility for maintaining and evolving the schema could be left to the community it serves. Individual data providers can then choose which schemas it will use, depending on the level of interoperability they can afford to support. For the architecture described above, the metadata used in search profiles is most relevant to data providers. Which profiles they support can be registered with the directory service; this information can be used by

gateways and agents for intelligently routing search queries. Furthermore, there is nothing barring a repository from supporting a schema that might be considered outside its discipline if it can at least in part be applied to its resources; doing so would further aid cross-discipline research. Finally, we can imagine smaller collaborations of repositories defining more detailed sub-schemas to serve special needs of the collaboration; if it conforms to the general framework, then it could be used by clients outside the collaboration.

We can imagine extending this concept to broader communities, such as to all of science or to the entire community of digital information users. A very good example of a top level metadata schema is represented by the Dublin Core. One of the nice features of the definition is that it is syntax- and protocol-independent. It provides a good starting point that communities can build onto “from underneath.” The W3C's Resource Description Framework (RDF) is expected to provide an approach to metadata definition that encourages interoperability across diverse applications. Use of XML will also encourage broader interoperability: the use of namespaces will allow one to mix schema together in a single query or response, and XSL allows for easy translation between schemas.

- The Vizier Catalogue Service at Strasbourg
<http://vizier.u-strasbg.fr/>
- Interoperable Systems for Archival Information Access (ISAIA)
<http://heasarc.gsfc.nasa.gov/isaia/>

5.3 Current Interoperability Efforts

The research community is starting to develop collaborations for the creation of these unifying infrastructures. The Thetis project is developing distributed information management technology to federate data collections of simulations and remote sensing data within specified disciplines. The technology builds upon metadata representations of data and simulation programs within a discipline, wraps data collections to support distributed queries search, retrieval and presentation of information via the Web. Data are published into the collections for use by other scientists. Further work builds upon ontological representations of knowledge within a discipline thus promotes data and program integration at the semantic level. Corresponding efforts within the US are represented by the Digital Library Initiative Phase II projects. In the DLI2 projects, one goal is support for interoperable services between disparate digital libraries. A European initiative to link catalogues of Earth Observation data is the INFEO system, that is based on the Catalogue Interoperability Protocol (CIP), which in turn is based on the Z39.50 protocol.

The research community is also developing information management technology through the Grid Forum interoperability effort. This community is promoting the ability to integrate remote data resources with remote computing capabilities. A similar effort is underway in Europe to build computational grid environments. Examples of these systems include the Globus and Legion distributed computing environments. **A major research goal** is the integration of the information management technology of the digital library community with the remote computation capabilities of the distributed computing environments. The combined system would then support data ingestion from remote sensors, publication into collections, analysis on compute platforms, comparison with simulations, and archiving for long term preservation. The steps used to manage the data flow can be represented by a workflow processing environment. The integration of workflow processing with information management and distributed computing can be an ultimate goal for a joint infrastructure development effort.

- Thetis
<http://www.ics.forth.gr/pleides/THETIS/thetis.html>
- INFEO - INformation on Earth Observation
<http://info.ceo.org/>
- Grid Forum Community
<http://www.gridforum.org/>
- Data Intensive Computing at the NSF NPACI collaboration
<http://www.npaci.edu/DICE/>

— Globus Project
<http://www.globus.org>

5.4 Security and authentication

In academia, authentication is treated as a sticky, difficult problem that is often left to the end of the project and then ignored; and yet a successful authentication framework is one that is built into the system from the beginning. Without authentication, we cannot rise above “toy data”, we cannot ingest, process, or deliver the data that real scientists are interested in, because this is generally not public. Clearly, in this age of hackers, security through obscurity is insufficient.

We should encourage projects that demonstrate easy to use, yet strong, authentication schemes, including ways to issue usage-permission to valid users, ways to log usage, ways to provide different levels of authentication. A most important facet of this problem is authentication in distributed systems, so that a user only needs to log-in once, yet multiple, heterogeneous services can be instantiated as a result. In other words, we need models by which one private service, accessed by a user of given authentication level, can access another private service at the same level. Current standards include the GSS (Generic Security Services API), and the GAA (Generic Authorization and Access API).

Access and control policies should be clear and unambiguous to those who create and use scientific data. The data may be public, or it may be restricted to a group of users listed explicitly or by domain name. The data may become public after a certain time, or there may be pricing policies. When the policies are clear, the implementation of them becomes possible. In particular, there should be agreement between EU and US agencies on these matters.

— Common Authentication Technology
<http://www.ietf.cnri.reston.va.us/html.charters/cat-charter.html>

6. Information Flow

It is generally much easier to accumulate data (sensors, experiments, simulations etc.) than to use it effectively. The goal must be to increase the productivity of the working scientist. This could be achieved both by giving individual scientists more intelligent and powerful tools with which to query and analyze the data and by helping to organize communities of scientists in their collective enterprises.

6.1 Design Patterns

Distributed systems offer a number of paradigms, or patterns, by which data, parameters, and code move through the system. The distinctions between these concepts themselves are vague — a large parameter set is simply data, and one machine’s program is just data to another. Some of the models are:

- Remote execution, where a program resident at another machine is caused to run there, perhaps using parameters as input,
- Query execution, where the remote program is a standard one (a “query engine”), and the input is a string of a particular language,
- A web server, where the output may cause either execution of a program or retrieval of a file,
- Retrieving a file, where there may be execution such as decompression or tape manipulation.

We see that these four traditionally distinct activities are all part of a continuum, the common thread being an initiating client and a responding, stateless, server: we see that there is little difference between a file, a query response, and a program output. Another pattern is based on agents, where requests, responses, and code are bundled into a travelling package, so that reduction of large amounts of data can occur close to where the data is stored. Many distributed systems are based on third-party transfer, where a client process informs a receiver that

data is coming, then requests a sender to transfer the data to the waiting receiver. These models can be complicated when the boundaries between data and program are blurred, and complicated even more when we consider scheduling and batch jobs, jobs that require data staging and temporary storage, and jobs that can be interactively steered as they execute.

Design patterns are ways to get perspective on a complex software system, ways to reuse design ideas as well as code; examples are the Model-View-Controller pattern and the Observer and Factory patterns. Many see design patterns as one of the most compelling ideas to emerge in object-oriented programming over the past few years. **We should encourage projects** that try to do the same for distributed database systems, ways of thinking about such systems, ways to designate responsibility among users and administrators, ways to define taxonomy, perhaps most importantly ways in which potential users can absorb complexity in a gradual way that puts them in control at all times.

6.2 E-commerce Software for Science

A design pattern for distributed databases is already familiar to online shoppers; discovery of the site, browsing and selection of products, authentication, payment and shipping options. Often the server authentication will be connected to user-specific state, for example, font preferences or a selection of previously-used shipping addresses. While such a design approach may be insufficient for many scientific databases (for example, there is no provision to specify data processing at the server), it might be sufficient for some projects, and a good start for others.

Besides the fact that users are already familiar with this pattern, there may be software already available to implement it. Business software, at best, is cheap (compared to a graduate student, or a supercomputer), well-documented, and robust. Unlike home-made software, new versions, with new features, appear regularly with no personal coding effort. All we need to do, for long-term projects, is to insulate ourselves from reliance on a single vendor by using open interfaces.

We should consider supporting academic projects that are closely partnered with industry, especially when they are on opposite sides of the Atlantic. The support is definitely not a subsidy to the business plan of the industrial partner, but rather should fund the insulation of the academic enterprise from the industrial partner! Specifically, **we should fund** the development of an open interface and the corresponding broker software, by which the two can effectively collaborate, and so that others can also join the enterprise.

6.3 Problem Solving Environments

In many regards, the needs and expectations of scientists are well ahead of online shoppers. The scientist will expect the capability to define a complex sequence of processing to be applied to the data, data-mining operations that reduce a vast bulk of data to small, concentrated, informative objects; also to specify where the data comes from, and where computing operations are executed. It should be possible to set up this specification on a thin, personal machine, but have it executed, or at least scheduled for execution, on chains of superdata and supercomputing resources. Environments for such distributed Problem Solving Environments are often based on directed acyclic graphs (DAGs) of modules, or on object-oriented scripting languages.

A researcher may wish to apply an analysis or presentation service for one database component against the data holdings of a different part of the federation. The development of interoperable data manipulation capabilities requires the integration of several approaches: in the commercial world, such services are implemented with CORBA, Java RMI and the Java EJB model; in the distributed computing environment, there are remote execution environments, in which data sets and procedures are moved to an appropriate compute platform. It should be possible to take a defined structure for a data set, map the data set to the structure required by an analysis procedure, and then apply the procedure. Given information models that describe how the attributes used to define

a data set are organized, it should then be possible to federate collections. The Common Component Architecture is an effort in this direction.

- US DOE Common Component Architecture Forum
<http://www.acl.lanl.gov/cca-forum/>
- Enterprise JavaBeans (EJB)
<http://ftp.java.sun.com/products/ejb/index.html>
- Scripting Languages
http://dir.yahoo.com/Computers_and_Internet/Programming_Languages/Scripting_Languages/

6.4 Discovery and Workflow

Scientific digital libraries and information services are now commonly used for scientific research. A few short years ago, a scientist could bookmark a few favorite URLs and have access to most of the data relevant to her work that was available in electronic form. Today, it is increasingly harder to answer the question, what data and information exists about...? The difficulty is not just in the sheer volume of data now available, but in the variety of data and in the number of repositories where they might be found, and the existence of catalogues. Thus, today more attention is being paid to issues of interoperability and federating resources.

Conventional resource discovery, at the level of individual resources, is not powerful enough to provide advanced, value-added services (e.g. environmental planning, forecasting). It becomes increasingly important to be able to identify entire groups of scientific resources that can be combined with each to support a given task. Users, or software agents executing on behalf of users, for example, may need access to data that have to be produced (at run time) via a series of complex data manipulations and computations using the available data sets and programs.

Hence, it is important to identify such combinations of scientific resources, and store such descriptions in a meta-programming information system. The approach is to use domain knowledge in describing tasks at the semantic level. In other words, domain-specific tasks for data production are defined in terms of graphs of ontological terms, connecting resources with each other. This information can be stored in a form that is appropriate for automated processing, retrieval and interactive browsing (e.g. XML schemas). A step further is to abstract relations among resources in a set of generic rules, stored in a knowledge base. Using the inference capability of knowledge bases, it then becomes possible to derive, on demand, different workflow scenarios depending on resource availability and restrictions posed by the user (e.g. range and accuracy of data to be produced, cost and duration of the computation).

Different semantic nets (rules) and thus workflows can be specified that cater to the needs of different user groups. Semantic interoperability is of utmost importance here; the information system must correctly interpret the user's information request and translate it to the descriptions of the distributed, heterogeneous repositories that provide the resources for satisfying this request. This can be achieved using federated mediation or metadata brokering techniques (Christophides et. al.)

Efficient and robust execution of workflows requires an appropriate runtime system. Apart from the conventional properties of workflow runtime systems, for the case of distributed scientific systems, data and program wrappers are additionally needed to hide the differences among local storage formats and execution platforms. Furthermore, appropriate access control schemes must be designed to ensure the protection of sensitive and expensive resources (data and programs). It should also be possible to run a given workflow in batch or interactive mode, and to provide execution control — stop, pause, start, abort and so on.

- V. Christophides, C. Houstis S. Lalis, H. Tsalapata, "Ontology-Driven Integration of Scientific Repositories" in Lecture Notes in Computer Science 1649, R. Pinter and S. Tsur (Eds.)

6.5 Distributed Data Mining

One part of the scientific process proceeds from experimental observation to theory. Algorithmic developments in data mining and related disciplines now mean that parts of this process can be automated. Instead of the scientist having to search large amounts of data for interesting patterns or relationships automatic methods can detect such patterns (associations, clusters etc.) and bring them to his/her attention. However performing such analyses over very large data sets is computationally expensive. This is an active area of high performance computing research and its application to large scientific databases would be appropriate. If the data is physically distributed and heterogeneous in nature, an extra complexity is added, but intelligent querying can offer a solution to the problems of analyzing such data sets. **We recommend** that research be supported into distributed data mining (theory, supporting infrastructure and algorithms) and its application to large distributed scientific databases.

Such a capability would permit a further raising of the level of scientific co-operation. Scientists could exchange results in the form of partial models that would be formal and computer processable. The various communities should be encouraged to develop the infrastructure to support distributed querying and model exchange over the information and data models previously discussed.

— Data Mining Research Group at Imperial College, London
<http://ruby.doc.ic.ac.uk/>

7. Preservation of Databases

For how long should scientific data be preserved? This depends on the value of the data into the future, and also on the cost of maintaining the data in a usable state. Sometimes, the content of a database can be reproduced at reasonable cost, in which case there is no need to worry about long-term preservation — for example if the database contains the results of simulations, and the software environment and a suitable machine are still available. In some circumstances experimental data can be reproduced, perhaps data from an accelerator or protein shape. On the other hand, the data may be irreproducible — the light-curve of a supernova or the ecological response to an environmental assault. The most common circumstance is that the data can, in principle, be reproduced, but it would be difficult and expensive to do so.

It may also be that the scientific value of the data can be completely extracted within a year or so of the data being created, in which case it can, in principle, be deleted. But the reality is that we cannot always be sure that there is nothing more: and in general the value of a scientific database increases as related data is added. But there is another reason, too: we believe that in the future it will be more and more common for raw data to be published in tandem with the peer-reviewed paper with the knowledge content of the database.

We recommend that creators of scientific databases should be encouraged to consider in advance how the database will be used in the future. Preservation description information should be associated with the digital objects being preserved so that

- the chain of custody and processing history are available,
- quality of the data is specified, perhaps defined by the nature of those who produced it,
- relationships to other digital objects are recognized,
- the digital objects can be unambiguously identified, and
- their information content is not altered in an undocumented manner.

To facilitate preservation, there is a need for a clear designation of what constitutes the boundaries of a given information object, and a need to resolve pointers within the data object to other data, such as local file names, directory paths, and dead (or dying) web hyperlinks.

Tertiary storage and cheap secondary storage are intrinsically unreliable. Tapes, which are good enough for

conventional back-up and restore applications are hardly adequate when the system starts to become much larger. The present CERN system with over 20,000 tape mounts per week and over 4 terabytes of data movement per day requires constant service by both CERN and manufacturer experts. The failures are several per day, mostly recovered with no data loss but at the expense of heavy and intense labor efforts. Given the niche nature of this tertiary storage market, technology is evolving very slowly. There may be software solutions to cope with unreliable hardware: the system of the future must stress resilience to hardware faults and scalability.

Data Rescue Research Center

We see a need for a “data rescue research center”, investigating the economics, and perhaps diminishing returns, of long-term storage. The center would consider migration to new media and what kinds of metadata and catalogs will survive the test of time. The center would consider, for example, the economics of:

- Refreshment - the replacement of a media instance with one of the same type,
- Replication - copying the full information content and the packaging information used to delimit this information content, to a new media instance of the same or different type,
- Repackaging - a copy where there is no change to the full information content, but some change to packaging information takes place,
- Transformation - a copy where there is some change to the full information content

— Reference Model for an Open Archival Information System
http://ssdoo.gsfc.nasa.gov/nost/isoas/ref_model.html

8. Education and Outreach

There are no guidelines as to what every domain scientist should know about IT. Hence we cannot, as a rule, reasonably expect them to substantially contribute during the definition of data representations or the creation of metadata. This is a serious deficiency as the modeling of scientific data should be the overall responsibility of the domain scientist (she or he knows best for what purpose the data are collected). **We recommend** investigation of ways to standardize requirements for IT-courses in the curricula of the domain sciences with emphasis in data modeling and use of databases. We should develop, implement and evaluate pilot-courses at the undergraduate and the graduate levels (extending it to include continuing education).

A sample syllabus for an undergraduate course could include topics such as: Information workplace and archival research, how information sources can be discovered, ways to transmit data, information models and formats, data processing and summarization, presentation of data, spreadsheets, visualization of simple and abstract datasets, dynamic and interactive visualization, data management, database systems, the relational model. Interdisciplinary work and examples are needed. This could be followed by graduate courses on managing scientific data. There would also be summer research programs to bring together undergraduates with domain scientists and also with computer scientists.

Such an educational initiative will almost immediately bring a better return on any investment in large scientific databases because scientists will be better equipped to handle their (expensive) data more responsibly, and also because scientists who choose a career in teaching can provide their students better access to scientific data and information.

- "Being Fluent with Information Technology". Committee on Information Technology Literacy, National Research Council, National Academy Press, Washington DC, 1999.
- "Scientific Database Management". Edited by J.C. French, A.K. Jones, J.L. Pfaltz. Computer Science Reports No. TR-90-21 (final report) and TR-90-22, University of Virginia, Dept. of Computer Science, Aug. 1990.
- Course on Software and Information Tools for Natural Scientists, ETH Zurich
<http://food.ethz.ch/info>

9. Funding Mechanisms

9.1 Recommendations for review process of funding

One possible model for proposal review is that both the EC and the NSF review each proposal separately, and only those which are approved by both will be funded. However, **we do not recommend** this model, as we feel it discourages innovation and collaboration.

On the contrary, **we recommend** that there should be a single review process and panel, with a single collaborative proposal, with copies going to both EU and US agencies. There should be discipline-specific as well as information-technology funding sources.

Once the joint EC/NSF panel has made a recommendation, **we recommend** that each proposal either be accepted by BOTH the EU and US agencies, or that it be rejected by both.

We recommend that the solicitation should bind the agencies to have proposals reviewed by a certain date, perhaps three months from the submission deadline.

To provide a strong coordination between projects funded under this solicitation, we recommend coordination of all activities funded under this program. In particular, there should be yearly PI meetings, and perhaps additional funding to one site for this purpose.

In the European Union, there is a predilection for more applied work than is funded at the US-NSF. We feel that these are complementary approaches, and do not create an impediment to collaboration.

9.2 Scale and Type of Projects

We first describe ways to fund the traditional research project — perhaps a few personnel on each side of the Atlantic, collaborating through shared data and code, through conferencing, through travel, sabbatical visits, and workshops. We also describe another approach, based on the US Science and Technology Center concept, a larger, longer-lived effort.

Traditional Research Teams

Awards might be made for supporting a specific research topic, where the complete strategy is defined in the proposal, and where substantial progress can be made with only a few people. Examples might be the federation of two existing scientific databases for a well-defined scientific purpose, the addition to a working database of a new information technology, or the use of data from one side of the Atlantic by a scientist on the other side.

Government funding agencies often require industrial participation or other cost-sharing before funding; however, in this case, we feel that this is not necessary. It is the application scientists, either in academia or industry, who are starting to accept the idea of scientific inquiry through databases, and we feel that if they are initial target partners, then there will be a strong leverage. In industry, however, there is little acceptance of the kind of scientific database that we are discussing here, (except in bioinformatics) and we feel that demanding industrial involvement would be an impediment to many projects that will not bear fruit in the commercial world for two or three years.

An Expedition Center for Large Scientific Databases

This workshop proposes the establishment of a virtual “center” (actually hosted at multiple geographical sites) similar in scope and thrust to the US-NSF Science and Technology Center concept. Essential features include:

- **unification** of information and knowledge management between US and Europe
- strong **leadership** and continuity of purpose,
- **funding** in the millions per year,
- **longevity** of 5 years or more,
- **flexibility** to seize new opportunities quickly and to shift the agenda rapidly.

The center would be a network of excellence in specific research domains, supporting both basic and applied research, with large-scale testbeds and large-scale demonstrations. It would have a strong education and outreach component. There might be two or three sites in the EU and US with independent funding for visitors, travel, and workshops. These regional centers could share the facilities of independently-funded facilities to create large-scale demonstrations and prototypes; such sharing could be achieved by collaboration agreements or by rental.

The expedition center would be a geographically distributed research institute, emphasizing trans-Atlantic teams, providing an intellectual home for a wide range of activities:

- Basic research as well as applied research and development,
- Testbeds for the integration and testing of technologies,
- Organization of workshops, summer-schools,
- Organization of specific areas of emphasis,
- Education, training, and other outreach,
- Coordination of other activities funded under the rest of the program,
- Liaison to other activities, for example the Framework 5 in the EU, the Computational Grid and Data Grid Forums in the US, and the Digital Library Initiatives in the US.

Why the Expedition Center Concept?

The work proposed in this report involves cross-cutting activities that work best if enough experts can work together: for example standards in data formats, information models, and protocols. We feel that the Expedition Center is a suitable vehicle to create and maintain critical mass so that these activities can be carried forward, carried into the scientific community, and eventually carried into industry. In addition, the Center can provide flexibility to respond quickly to technology changes and to encourage blossoming projects.

Management of the Expedition Center

Proposals to set up the Expedition Center should specify management quite closely, showing regard for both the EU and the US funding mechanisms. However, we do not believe that the Expedition Center should be the repository of record for the data it holds, as that would reduce flexibility and discourage new technology.

- NSF Science and Technology Centers
<http://www.nsf.gov/od/oia/programs/stc/>
- President's Information Technology Advisory Committee Report (PITAC)
<http://www.ccic.gov/ac/interim/>

10. Participants of the Workshop

- Peter Allan, Rutherford Appleton Laboratory, p.m.allan@rl.ac.uk
- Giovanni Aloisio, University of Lecce, aloisio@sara.unile.it
- Alberto Apostolico, Purdue University, axa@cs.purdue.edu
- Dominique Boutigny, LAPP - CNRS/IN2P3, boutigny@lapp.in2p3.fr
- Peter Brezany, University of Vienna, brezany@par.univie.ac.at
- Julian Bunn, Caltech and CERN, julian@cacr.caltech.edu
- John Darlington, Imperial College, jd@doc.ic.ac.uk
- James Frew, University of California, frew@bren.ucsb.edu
- Fabrizio Gagliardi, CERN, Fabrizio.Gagliardi@cern.ch
- Hans Hinterberger, ETH Zurich, hinterbe@inf.ethz.ch
- Catherine Houstis, University of Crete, houstis@ics.forth.gr
- David Maier, Oregon Graduate Institute, maier@cse.ogi.edu
- Paul Messina, Caltech and US DOE, messina@cacr.caltech.edu
- Reagan Moore, San Diego Supercomputer Center, moore@sdsc.edu
- Francois Ochsenbein, Observatoire de Strasbourg, francois@astro.u-strasbg.fr
- Raymond Plante, NCSA, University of Illinois, rplante@ncsa.uiuc.edu
- Daniel Reed, University of Illinois, reed@cs.uiuc.edu
- Donald Sawyer, NASA, donald.sawyer@gssc.nasa.gov
- Gunter Schreier, German Aerospace Center, Gunter.Schreier@dlr.de
- Arno Siebes, CWI, Amsterdam, arno@cw.nl
- Alexander Szalay, Johns Hopkins University, szalay@jhu.edu
- Roy Williams, Caltech, roy@cacr.caltech.edu
- Maria Zemankova, National Science Foundation, mzemanko@nsf.gov (observer)