

Proceedings

PV-2004

**"Ensuring the Long-Term
Preservation and Adding Value to
the Scientific and Technical Data"**

5-7 October 2004

ESA/ESRIN, Frascati, Italy

Programme Committee

Chairman

Gian Maria Pinna - ESA, Frascati, Italy

Vice-chairmen

Vincenzo Beruti - ESA, Frascati, Italy

Philippe Mognaud - ESA, Frascati, Italy

Members

Lars Alm - SSC, Kiruna, Sweden

John Faundeen - USGS, Sioux Falls, US

Francoise Genova - Centre de Données Astronomique de Strasbourg, Strasbourg, France

Paolo Giommi - ASI, Rome, Italy

Bruno Greco - ESA, Frascati, Italy

Claude Huc - CNES, Toulouse, France

Tom Lankester - Infoterra, Farnborough, UK

Patrick Mazal - CNES, Toulouse, France

Eberhard Mikusch - DLR, Oberpfaffenhofen, Germany

Nestor Peccia - ESA, Darmstadt, Germany

Gilles Petitjean - EADS/S&DE, Paris, France

Nicolas Picot - CNES, Toulouse, France

Giancarlo Pittella - ESA, Frascati, Italy

Donald Sawyer - NASA/GSFC, Washington, US

Workshop secretariat

ESA Conference Bureau

P.O. Box 299

2200 AG Noordwijk

The Netherlands

t: +31 71 565 5005

f: +31 71 565 5658

e: esa.conference.bureau@esa.int

Table of Contents	Page
Session Technology And Standards	
From the OAIS Reference Model to a Data Center Organizational Model <i>C. Huc, CNES</i>	9
Object Query Language - Enabling Service for Earth Observation Product Processing, Access and Dissemination <i>S. Kiemle, A-K. Schroeder-Lanz, C. Reck, DLR</i>	17
Developing a System to Support Cross-Instrument and Cross-Mission Search and Retrieval of Planetary Data <i>T. C. Stein, E. A. Guinness, Washington University</i>	25
The Planetary Science Archive Introduction And Overview <i>J. J. Zender*, G. Schwehm**, C. Arviset**, **ESA/ESTEC, **ESA/ESAC</i>	31
Storhouse: an Affordable Strategy for High- Volume Digital Preservation <i>A. Dyball*, D. Clements**, *Filetek UK Ltd, **Filetek Inc.</i>	39
Archiving and Management of Geospatial Data <i>T. Ohlhof, R. Syffus, ESG Elektroniksystem Und Logistik Gmbh</i>	47
“The High Speed ESA Earth Observation Network Project for Improving Data Exchange And Access to the User” <i>M. Ciccarelli*, M. Forcada Arregui**, Datamat S.P.A*, ** ESA/ESRIN</i>	55
Virtual Access to Information:an Emerging Concept <i>J-P. Antikidis, CNES</i>	63
Operational Solutions For Long-Term Preservation and Advanced Access Services <i>G. Petitjean, EADS DCS</i>	69
EUMETSAT Unified Archive and Retrieval Facility: Status and Plans <i>F. Cade, EUMETSAT</i>	77
Understanding Indexing-Contextures With a View to Add Value on Models and Datasets in E-Science <i>Y. Liu*, F. Yang*, A. Coleman**, *University Of Aberdeen, **University Of Arizona</i>	85
Session Added Value Services	
Use of Browse or Preview Images in Support of Data Access <i>S. Doescher, R. Sunne, M. Neiers, USGS/EROS Data Center</i>	95
SITOOOLS:Enhanced Use of Laboratory Services and Data <i>T. Levoir*, L. Carbonnaux**, C. Gougeard**, *CNES, ** CAPGEMINI</i>	103

SIPAD-NG: a Multi-Field System for Accessing Scientific Data and Added-Value Services <i>D. Heulet*, E. Houdet**, * CNES, **CS SI</i>	111
Data Description Based on XML Dictionaries <i>A. Lucas, D. Minguillon, CNES</i>	119
Infrastructures to Improve User Access to EO Data <i>P. Seifert*, S. Kiemle*, E. Mikusch*, W. Wildegger*, M. Boettcher**, H-J. Wolf*, H. Maass*, K-D Missling*, * DLR, ** Werum Software & Systems</i>	127
Integrating Research Data into the Publication Workflow: the Ebank UK Experience <i>R. Heery*, M. Duke*, M. Day*, L. Lyon*, M. B. Hursthouse**, J. G. Frey**, S. J. Coles**, C. Gutteridge**, L. A. Carr**, * University Of Bath, **University Of Southampton</i>	135
Added-Value Services in Ether Atmospheric Chemistry Data Centre <i>C.Boonne*, F. Girod**, *IPSL, **CNES</i>	143
Information Mining and Understanding of Data Massive: a New Era of Methodology and Technology <i>Mihai Datcu, DLR (German Aerospace Centre)</i>	149
Semantic Web Technologies for Searching, Retrieving and Adding Value to Large Scale Scientific Data <i>C. Houstis, M. Pitikakis, G. V. Vasilakis, M. Vavalis, Institute Of Informatics And Telematics</i>	151
ESA RSSD Science Archives User Interfaces and Inter-Operability Systems <i>C. Arviset, J. Dowson, J. Hernández, I.Ortiz, P. Osuna, J. Salgado, G. San Miguel, A. Venet, ESA/VILSPA</i>	159
Discover, Browse and Access Ocean Data On Internet <i>S. Baudel*, F. Blanc*, T. Jolibois**, V. Rosmorduc*, * CLS, ** COFRAMI</i>	165
On Using Grid Computing Technology with the OAIS Reference Model to Provide Persistent Archive Environments and Value-Enhancement for Nasa Earth Science Data <i>B. R. Barkstrom, D. E. Corder, NASA Langley Research Center</i>	171
The ESA Service Support Environment - Exploitation of the Long Term Archives <i>S. D'Elia, P. Marchetti, ESA/ESRIN</i>	179
Session Lesson Learnt	
Value Added Services for Mars Express' ASPERA- 3Data System <i>J. Mukherjee, M. Muller, A. Galus, S. Jeffers, C. Gonzales, J. Pardue, Swri</i>	187
From Data to Knowledge in Earth Science, Planetary Science, and Astronomy <i>E. R.Dobinson, J. C. Jacob, T. P. Yunck, Jet Propulsion Laboratory</i>	193

U.S. Geological Survey Scientific Records Appraisal Tool <i>J. L. Faundeen, C. J. Wippich, U.S. Geological Survey</i>	201
Twelve Years of User Services for Ocean Topography Users: Aviso Experience and Lessons Learned <i>V. Rosmorduc*, N. Picot**, *CLS, ** CNES</i>	209
Preserving Access to Legacy Information Through Data Migrations at NSSDC: Experiences and Lessons Learned <i>D. Sawyer, H. Kent Hills, P. Mccaslin, National Space Data Center</i>	215
The Lifecycle Of NASA's Earth Science Enterprise Data Resources <i>K. Mcdonald*, R. A. Mckinney**, T. B. Smith**, R. Rank***</i> <i>*NASA, **USGS/EROS, ***NOAA/NESDIS</i>	223
Session Future Prospects	
Digital Library: Improving The Accessibility of the Russian Satellite Data in Support of the Environmental Monitoring <i>Efim Kudashev, Space Research Institute Of Russian Academy Of Sciences</i>	233
Providing Authentic Long-Term Archival Access to Complex Relational Data <i>S. Heuscher, S. Jaermann, P. Keller-Maxer, F. Moehle, Swiss Federal Archives</i>	241
A Planetary Data System for the 2006 Mars Reconnaissance Orbiter Era <i>J.S. Hughes, D. Crichton, G. Crichton, R. Joyner, S. Kelly, C. Mattmann, J. Wilf, Jet Propulsion Laboratory</i>	263
The UK Digital Curation Centre <i>D. Giarretta*, P. Burnhill**, S. Ross***, L. Lyon****, P. Buneman**, * CCLRC, ** University Of Edinburgh, ***HATII, ****UKOLN</i>	269
Digital Libraries And Grid Technologies as Infrastructure for Earth Observation Data Archives Exploitation an Long- Term Preservation <i>L. Fusco*, J. Van Bemmelen*, V. Guidetti*, D. Castelli**, * ESA/ESRIN, **CNR</i>	277
Information Mining in Remote Sensing Images - the KIM, KES and KIMV Projects <i>A. Colapicchioni*, S. D' Elia**, M. Datcu***, *Advanced Computer Systems S.P.A, **ESA/ESRIN, ***DLR.</i>	285
SPASE - Space Physics Archive Search and Extract <i>C. Harvey*, J. Thieman**, T. King***, D. Aaron Roberts**</i> <i>* CESR, **NASA, ***UCLA,</i>	291

Session Technology and Standards

From the OAIS Reference Model to a Data Center Organizational Model

Claude HUC

*Centre National d'Etudes Spatiales
18 avenue Edouard Belin
31401 Toulouse Cedex 9, France
claude.huc@cnes.fr*

1. INTRODUCTION

The OAIS [1] Reference Model is both a CCSDS recommendation and an ISO standard. The latter, which is broad in scope and well reasoned, provides an insight into the complexity and scope of the long-term archival of digital information. It lays down the concepts and vocabulary that are needed to understand the problem. The OAIS model is currently an essential basis, recognized as such by members of the scientific, industrial and heritage communities.

We have come a long way in understanding the issue and its multiple facets—whether relating to organization, technology, legal aspects or standards. However, understanding the issue is not enough: **the real need is to produce concrete, applicable solutions**. Because one of the major causes of the difficulties involved in ensuring the long-term future of digital data is the constant obsolescence of technologies, we sometimes tend to focus our analyses and solutions on the technical side of the problem. Though we should not underestimate the importance of technical issues, it appears clear today that solutions will only be developed when several distinct factors are combined. These factors include suitable organization, competent people at each step, a set of applicable procedures and standards, appropriate technical resources and a certification procedure for archives that respect the standards.

This article focuses on the organizational aspects that appear so critical. The organizational model we propose is the result of considerable brainstorming that includes considerations such as the global vision of the OAIS model, feedback on the practical archiving of data and scientific documents at CNES, the current structuring of technologies into distinct sectors and the human factors related to the wide range of skills required.

The organizational model described below is not the only valid model. It is, however, a model that works. It will no doubt change in the future to keep abreast of the new jobs and technologies that will emerge and of which we currently know nothing. It will also be developed to take into consideration an analysis of the successes and failures of practices implemented.

2. MODEL STRUCTURE

Digital data archival is so complex and requires such a wide variety of competencies that it cannot be resolved as a single unit. We therefore broke the initial problem down into much smaller, independent problems. Our organizational considerations began with the OAIS functional model shown in figure 1. All the functions identified by this model also need to be covered by an organizational model.

The four vital operational functions of the OAIS Model are ingest, archival storage, data management and access. Practical experience had already shown that it was possible to split these operational functions up into three separate services:

- an 'Ingest' service in charge of all ingest-related tasks,
- an 'Archival Storage' service corresponding closely to the archival storage function,
- a 'Data Management and Access' service to manage and offer users access to archived data and metadata. The 'Data management' and 'Access' functions are closely linked in that the information search criteria offered to users by 'Access' are just the concrete expression via a human/machine interface of a data management search system .




Figure 2: The archive organizational model

3. THE INGEST SERVICE (IS)

3.1 Functions and Responsibilities

The 'Ingest Service' (IS) is responsible for collecting data objects from producers and for all the tasks involved in transforming the objects handed over by producers into digital objects that have all the qualities needed for them to be archived.

The IS is primarily responsible for interactions between the archive and data producers. Producers may be other services within a space agency, research laboratory, other scientific or technical agencies or institutions etc. For each producer, the various digital data objects to be transferred to the archive must be defined (in a document known as the archival plan or data management plan). The transfer process must then be implemented and the objects received by the archive must comply with expectations. A global view of this process is given in the CCSDS 'Producer Archive Interface Abstract Standard' [2]. It is important to note that generally the different producers are not all capable of implementing archive rules on data and metadata formats. This is usually the subject of negotiations leading to an agreement acceptable by both parties. At this point, discussions are not only on identifying the data and documents to be archived but on the exact format in which these documents will be transferred, the type and form of metadata that could accompany these documents and the way in which these digital objects will be transferred to the archive.

Furthermore, the archive will often have to collect information from several different producers to constitute an intelligible archival information package: in the case of a science experiment aboard a satellite the archive shall, for example:

- collect data in a form calibrated by the PI's team. There should be a semantic and syntactic description of collected data. It must also be sufficiently documented and accompanied by metadata in a standardized form,
- collect general information on the space mission from the space agency that conducted the mission,
- collect data on orbit, attitude etc. from the Mission Control Centre,
- etc.

The acceptable or recommended data formats for long-term data preservation must satisfy precise criteria. They must be standardized and independent of the software used to create them. It must also be possible to describe (syntax and semantics) the data objects exhaustively. The Ingest Service must therefore:

- receive objects from producers and check their compliance with the established plan,
- change the data and metadata format when necessary (files delivered in MS Word format, for example, may be transformed into PDF/Archive format, and text files containing metadata may be transformed into structured XML files),

- assign the digital objects received a unique identifier consistent with the archive's naming space,
- add to metadata by placing received objects in a contextual relationship with other archived objects or with documents available in other archives,
- transfer all archivable data objects (data and metadata) to the Archival Storage Service,
- transfer metadata and any objects that are to be accessible on line to the Data Management and Access Service.

The OAIS model insists on the need to gather and organize information such that it remains perfectly understandable both today and tomorrow by the user community. This is no trivial issue, but lies at the heart of the Ingest Service's responsibility.

3.2 External Interfaces

- Interfaces with Producers: the way that data objects are to be transferred and that communication is to take place between the producer and the archive must be specified in a document approved by both parties.
- Interfaces with the Archival Storage Service. These are very simple interfaces. They may be resumed by a few actions that may be implemented from a workstation within the Ingest Service. A realistic example of actions enabling the IS to transfer a digital object to the AS service follows:
 - ⇒ *Log-in to the AS service (initiated by the IS), authentication*
 - ⇒ *Request to store a digital object, specifying the object's identifier and the type of service expected (cf. the archival storage section)*
 - ⇒ *File transfer*
 - ⇒ *The AS service sends an acknowledgement*
 - ⇒ *Session closed*

The process implemented in this way may be manual—implemented by an operator on a case-by-case basis depending on needs—or automatic, if many data objects are to be transferred. The actions defined above can obviously be formulated and organized differently depending on the constraints involved and choices made. It is essential for this interface to be independent of activities upstream of the Ingest Service and to avoid prejudicing activities carried out downstream by the AS service. These same principles apply to interfaces with the DMA service.

- Interfaces with the Data Management and Access Service: the DMA receives standardized metadata files of various levels. They may include descriptions of collections and subcollections or descriptions and the identification of single data objects. The Ingest Service may also transfer specific digital objects that are useful when seeking information in the archive (e.g. diagrams or graphs).

The AS service has no knowledge of the content of the files it receives. It receives a bit stream accompanied by its identifier and a storage instruction and it must guarantee the future of these bits. The DMA service does, however, process the information received and organizes it within a database so that it is accessible to users. The resulting interface between the Ingest and DMA services is necessarily more complex even if it is standardized. CNES applies ISO standards on data dictionaries and their XML representation [3] and [4] to these interfaces.

3.3 Technical facilities required

The hardware, software and communication means needed to receive digital data objects sent by controlling agencies do not have any specific characteristics although it must be possible to decide on a case-by-case basis to secure transfers to as to authenticate received objects and guarantee their integrity to the sender. Facilities must be geared to the volume of data handled and frequency of transfers. A software suite will be required to help prepare data and metadata.

3.4 Skills required

Several kinds of skill that are today found separately need to be combined within the Ingest Service:

The skills of a archivist capable of defining, in liaison with the producer, which information needs to be archived and of organizing this information within a structured unit,

The skills of a computer scientist specializing in data management and the representation of information in digital form, in order to define the data and metadata formats suitable for long-term preservation, to check that

delivered objects comply with these formats, to change format if necessary, to specify the development of computer tools needed by this service, to develop and then operate them. Skills like these focusing on digital representations also assume a general knowledge of computer science. Both skills are combined in a new profession described here as a digital data manager or digital information archivist.

To these two skills must be added a third one: *the skill of a scientist* capable of checking the intelligibility and completeness of archived data.

3.5 Feedback from CNES's own experience

From a practical viewpoint, collecting a complete, organized and suitably described set of digital objects is the most difficult task and also the most costly in terms of human resources.

4. THE ARCHIVAL STORAGE SERVICE

4.1 Functions and Responsibilities

The AS service is responsible for the long-term preservation and integrity of digital objects in its charge. The AS service is unaware of the content of the objects it receives but must guarantee their long-term future whatever the media on which they are stored.

The Archival Storage Service has two clients, the Ingest Service and the DMA service.

The IS transfers digital objects to the AS service for preservation. The IS may request a particular class of service for these objects. Service classes specify the services expected of the AS service in functional terms, not technological terms. They may specify the retrieval time required (immediately, within one hour of a retrieval request, or within normal office hours, for example) or the degree to which preservation is vital. If an archive has a copy of digital data in another archive (for reasons of convenience for its users), it may not be necessary to have two or three backups for these objects, which it would be if the archive was fully responsible for the data objects in question. The DMA service sends object retrieval requests to the AS service depending on its own needs and user requests.

The Archival Storage service is responsible for all the activities needed to ensure the continued integrity of digital objects:

- storage of objects on media together with one or more backup copies kept in a separate place,
- constantly monitoring the condition of media used (number of write operations to each medium, bit error rate etc.),
- periodic replacement of less reliable media by new media,
- migration to new media more appropriate to its activities according to storage technology development (migrations may be periodic or continuous depending on the policy applied).

4.2 External interfaces

Interfaces with the Ingest Service have already been mentioned. Interfaces with the DMA service may take the simple form of the following example:

- ⇒ *Log-in to the AS service initiated by the DMA service, authentication*
- ⇒ *Request for the retrieval of a data object defined by its identifier*
- ⇒ *The AS service transfers the file to the DMA service*
- ⇒ *Session closed*

It is vital that interfaces with the AS service remain stable over time, whatever the developments or migrations within the service. Stability is the basis for independence between the three services. No major change within any one service must be allowed to have an impact on the other services.

4.3 Technical facilities required

The AS service must have all the data storage equipment needed to fulfil its role: read-write devices, storage media, automated equipment and software needed to implement all operations. Depending on needs, it may operate 24 hours a day or within certain times on certain days.

A number of manufacturers offer systems that satisfy AS service needs to a greater or lesser extent. However, we should not forget that we should never fully trust the performance of systems or storage media available on the market. The archive remains fully responsible for data preservation.

4.4 Skills required

The AS service requires computer scientists specialized in:

- Managing large sets of stored files duplicated on various kinds of storage media,
- High-speed network technologies for communication with service “clients”,
- High-capacity storage technologies, storage automation, storage media, their characteristics and reliability,
- The assets used to monitor the state of media and their implementation,
- The ability to keep a system open 24 hours a day in operational order and develop the system to keep abreast of developments in technology and volume.

4.5 Feedback from CNES’s own experience

An Archival Storage service known as STAF (*Service de Transfert et d’Archivage de Fichiers*) was set up at CNES in 1995. It stores several hundred terabytes of data. After ten years of operation, the feedback is extremely positive on numerous aspects. The service is used for the long-term preservation of data from spaceborne experiments. It is also used to archive the agency’s technical documents and, finally, by numerous projects for their own short- or long-term storage purposes. Its integration in CNES’s in-house information system makes it a unique and indispensable tool.

Its interfaces with “client” entities have remained stable since its creation.

We should emphasize that the need for specialist skills in storage technologies means that such a service is only viable (in terms of global cost) beyond a certain critical mass. We have not said that each archive should have an AS service. Other organizational options include:

- A joint Archival Storage service common to several different archiving centres,
- An AS service shared by archives and other services within the same organization (the case at CNES),
- An AS service in the form of an independent service provider.

5. THE ‘DATA MANAGEMENT AND ACCESS’ SERVICE (DMA)

5.1 Functions and responsibilities

The DMA service is responsible for managing the wealth of information preserved by the archive, and its communication to authorized users.

It must set up and maintain a computer system providing constant access to complete and detailed knowledge of the archived objects and their collection within a set of hierarchically arranged collections (archive groups and subgroups). It must expand and enrich its data base using metadata provided by the Data Management and Access service. This system must also grant users remote access via a user-friendly interface which allows them to:

- Know the content of the communicable part of the archive,
- Seek data objects of use to their own work (applying selection criteria based on metadata, for example),
- Select, order and retrieve data objects of interest ,
- Transform archived data objects if necessary before providing them to a user (e.g. changing format or adding value).

The search for useful data can be based not only on metadata but other complementary techniques such as browsing or data mining. Data objects may be retrieved either via a network or by copying the object onto a commonly used medium such as CD-ROM, DVD or DLT depending on the volume.

The DMA service is also responsible for managing relations with the user community, including authorizations, access rights and invoicing if applicable.

5.2 External interfaces

Interfaces with the AS and IS have already been discussed.

As the archives in question are digital and considering current technological developments, the interface for archive users should enable users to search for and order data objects remotely over a network via a suitable graphical interface. The confidentiality and sensitive nature of stored information may lead to restrictions on general viewing.

5.3 Technical facilities required

The system set up by the DMA service is underpinned by database and Internet communication technologies in the main. Systems that partly or fully satisfy DMA service needs are or will soon be available on the market, which will limit the cost of dedicated software development. The DMA service should be able to copy digital objects on media for distribution. In certain cases—independent of the AS service—the DMA service may need to store data objects which have to be available immediately on line for user, thus requiring a storage capability (usually on a disk).

5.4 Skills required

Two kinds of skills are required. Firstly, the skills of a computer scientist specializing in:

- Modeling and information search processes,
- Database technologies,
- Internet technologies and languages (graphic interface for navigators etc.),
- The operational maintenance of systems serving large or even vast user communities.

Secondly, general knowledge of the subject and issues involved in archiving:

- Knowledge of the categories of data handled,
- Knowledge of metadata and data selection criteria geared to user needs.

5.5 Feedback from CNES's own experience

Data management and communication capabilities have been implemented to provide access to space science data in different fields (such as astronomy and oceanography). Despite the diversity of digital objects and rationales peculiar to each field, the challenge (which has already been taken up) is to reduce costs by the use of generic systems able to be tailored to each discipline. This is the case for SIPAD-NG and SITOOLS which will be presented during the PV2004 Workshop.

6. COORDINATION

A coordinating structure is needed to guide the archive. This is a management function. The **Coordinator** is the director and as such truly responsible for data archival as per the responsibilities defined in the OAIS model.

The coordinator's role is to precisely define which service does what and ensure the clarity of interfaces between services. They must be managed so that they cannot change without having to assess the impact of these changes on the other services. The coordinator also organizes activities relating to elements in common: the information model, dictionary of deliverable data objects and the identification of objects. The information model, for example, gives an overall, ordered and coherent view of the information objects managed by the archive. This global view includes object identifiers, their relationship and their assignment to collections....The IS must know this model because it is the IS that assigns identifiers to the objects and defines the place for these objects within the model and within the DMA service, which organizes the model in the form of a database. The coordinator also has to approve the standards and procedures (preferably standardized themselves) applied within each service.

Conclusion

We are all aware that technology will continue to develop, with new operating systems replacing previous ones and new programming languages developing (15 years ago there was no WWW, Java or Web services!) while other languages will disappear. **Information will remain.** We do not know the nature or size of technology leaps to come and the changes in use that will follow on. **This is why we have knowingly followed paths based on knowledge of information structure, syntax and semantics rather than attempting to maintain one technology in the long term.** Experience has shown that to keep software up to date in the long term is much more complicated and finally more costly than preserving information in a form independent of technology.

The organizational model proposed is based on this choice. It is intended to **foster the advent of concrete, applicable solutions.** It is also based on an analysis of skills and professions, and on feedback from CNES's own experience which has confirmed our rationale. Such an organization must be open to external inspections and audits. The digital archive must be able to demonstrate—through its organization, its resources, its teams and applicable standards and procedures—that it is capable of fulfilling its mission, the long-term preservation of the digital information for which it is responsible. This raises the issue of the 'Certification' of digital archives.

REFERENCES

- [1] CCSDS 650.0-B-1., Reference Model for an Open Archival Information System (OAIS). *CCSDS Blue Book. Issue 1. January 2002, ISO 14721:2003*, <http://www.ccsds.org/CCSDS/recommandreports.html>
- [2] CCSDS 651.0-R-1., Producer-Archive Interface Methodology Abstract Standard (PAIMAS). *CCSDS Blue Book Issue 1, May 2004*. <http://www.ccsds.org/CCSDS/recommandreports.html>
- [3] CCSDS 647.1-B-1., *Data Entity Dictionary Specification Language (DEDSL) - Abstract Syntax (CCSD0011)*. *CCSDS Blue Book. Issue 1. June 2001 ISO 21961:2003*, <http://www.ccsds.org/CCSDS/recommandreports.html>
- [4] CCSDS 647.3-B-1 Data Entity Dictionary Specification Language (DEDSL)—XML/DTD Syntax (CCSD0013). *CCSDS Blue Book. Issue 1. January 2002, ISO 22643:2003*, <http://www.ccsds.org/CCSDS/recommandreports.html>

OBJECT QUERY LANGUAGE - ENABLING SERVICE FOR EARTH OBSERVATION PRODUCT PROCESSING, ACCESS AND DISSEMINATION

Stephan Kiemle, Anne-Katrin Schroeder-Lanz, Christoph Reck

German Aerospace Center (DLR)

German Remote Sensing Data Center (DFD)

Oberpfaffenhofen, 82234 Weßling, Germany

Email: Stephan.Kiemle@dlr.de, Anne-Katrin.Schroeder-Lanz@dlr.de, Christoph.Reck@dlr.de

INTRODUCTION

Multi-mission product management systems have an uncontested primary focus on long term data preservation. However, sophisticated archive systems are incomplete if they provide only monosyllabic product identification and retrieval functions. The increasing diversity and amount of products and an increasing complexity of applications require flexible and efficient query capabilities supporting spatio-temporal conditions, combination of related products and specific product metadata conditions and aggregations.

This paper presents the Object Query Language (OQL) provided by the Product Library of the Data Information and Management System (DIMS), the multi-mission ground segment infrastructure implemented and operated at DLR DFD. This comprehensive query language is the basis for an integrated inventory service supporting dynamic object-oriented metadata configuration, ad-hoc queries, event triggering, incremental query of huge result sets and full metadata inspection. OQL supports complex query conditions including object reference navigation and different spatio-temporal operators. The paper presents prominent examples of operational systems illustrating the advantages of OQL, enabling new possibilities for processing, access, subscription and dissemination of products.

The Product Library is the central DIMS subsystem responsible for product management and long term preservation [1], [2]. At DFD's facility it manages (mid 2004) 2 million product items, 50 TByte in 3 million files, of various earth observation missions such as SRTM, ERS, ENVISAT, MODIS, Meteosat, MSG, BIRD, GRACE, CHAMP, SRL and airborne imaging. The Product Library serves multiple processing chains with raw to level 3 products and enables data-driven workflows as systematic processing, subscription ordering and automatic upload to user service systems (Fig. 1). The Object Query Language interface of the Product Library is of primary importance for these processing, access and dissemination services supported by DIMS.




Fig. 1. The DIMS Product Library in its multi-mission infrastructure context

In the first section we will focus on the relevance of metadata in the multi-mission infrastructure and motivate the necessity of a powerful query language. We then introduce object-oriented metadata modelling and the basics of the

Object Query Language, and expose how a metadata model is used to support metadata standards. Our focus in this paper is the role of OQL in the multi-mission ground segment and we describe in detail different concrete scenarios taking large benefit of OQL. A section about query performance issues and limitations concludes this discussion.

METADATA – NUTS AND BOLTS FOR PRODUCT MANAGEMENT

Earth observation ground segments have to deal with digital data of very large size. In order to be able to identify and handle data products, characterising and describing metadata is stored with the data. The role of metadata evolved a lot since the very beginning of earth observation from space, where metadata was cryptically coded and often only available in the name of the product file. The first operational missions made dedicated archive systems indispensable, and metadata catalogue systems came up enabling efficient search and retrieval. Nowadays multi-mission ground segment infrastructures supporting end-to-end business processes through production, online archiving, user access, ordering and delivery impose even more requirements on data management in general and on metadata in particular.

Primary product characteristics are no longer the exclusive source of metadata (Tab. 1); metadata has also to reflect the production process, e.g. the bundling of single scenes, quality assurance results and the reference to input and auxiliary data. Furthermore, metadata has to serve multiple purposes in the multi-mission infrastructure, e.g. the selection of processing input products, the selection of order options, the monitoring of product-oriented workflows and the product presentation to end users. Metadata has to be compliant to international standards for interoperability of structures [3], [4] and services [5].

Table 1. Product metadata and its typical use

Metadata source	Typical parameters	Typical use cases
Platform (up to level 2)	Mission name Orbit/flight characteristics	Input selection for higher level processing
Instrument (up to level 2)	Viewing geometry Scan characteristics, sensor mode	Algorithm setup for higher level processing, Selection of order options
Reception (up to level 2)	Acquisition station/ingestion referenc	Acquisition, reception and processing delays analysis
Processing algorithm	Spatial and temporal coverage and resolution Species/spectral bands/layers Projection/geocoding properties Scene properties (e.g. cloud/land/sea coverage) Algorithm version, configuration version	Product search/presentation in user services, Selection of processing options, Tailoring of value-added processing
Processing environment	Time of processing Input and auxiliary data version, orbit quality Product revision, mission project phase	Version control for reprocessing, Criteria for long-term preservation
Quality assurance	Rate of missing and degraded records Rate of contents threshold violations Mean, median and variance of contents Explicit quality information	Product and instrument degradation analysis, Warranty compliance after ordering and delivery
Product management	Time of archiving, archiving expiration time Size and format Access permissions	Regular delivery to subscribed customers, Purging of expired products, Archive population analysis

Finally, metadata is essential for organising the long-term preservation of a large number of digital products. Archives are based on hierarchical storage management file systems and require balanced, human readable directory structures (composed of metadata values) which have to be configurable and mission-specific. Operators have to compute usage and population statistics taking into account specific product characteristics. Consistency analysis has to be driven based on metadata to ensure completeness and correctness of the acquisition and production process.

The definition of product metadata - product modelling - is a complex task and requires experience. In order to fulfil all requirements on metadata in the multi-mission infrastructure the operator designing a product model has to gather information from the mission operations, the instrument development and the product generation subsystems and has to interview the user community and respect the international metadata standards. Therefore data models have to be individually configurable, well documented and extensible. Once complex product metadata models are defined and used, a powerful query language has to enable its exploitation, allowing individual selections and complex conditions.

The DIMS Product Library provides tools for definition and configuration of data models. The data model can be extended with product types for new missions during operations without service shutdown. The Product Library query interface allows browsing the collections of individual product types and placing ad-hoc OQL queries on specific product parameters. Even though the metadata model design is object-oriented, the final storage of metadata is relational. The inventory middleware of the Product Library decouples application-oriented and comprehensive models from normalised and complex relational database structures, and automatically maps OQL into efficient SQL.

ABOUT METADATA MODELLING AND OBJECT QUERIES

Object-oriented Metadata Modelling

As explained above metadata serves multiple purposes and reflects various properties of product-related processes. A simple list of key-value pairs cannot satisfy these needs. The object-oriented modelling method (Unified Modelling Language UML) proposed by Booch and Rumbaugh [6] is suitable to map the world of earth observation data products. The UML static model for object structures used in the Product Library allows using the following modelling elements:

- Attributes of primitive data type (including valids, constants, constraints)
- Structured attributes (compositions)
- Collection attributes (sets and lists of multiple values or structures)
- Object types (instantiable item definition, e.g. collection type, product type, component type)
- Associations between object types
- Object type inheritance

Attributes, structures and collection attributes are used to model product metadata parameters. The object types define the name of the type, a list of attributes uniquely identifying instances and a set of attributes describing instances. Associations are used to relate e.g. collections, products and components. A product usually associates different components, e.g. a browse data component and a primary data component. A product may also associate predecessor products, the inputs that were used to generate the product. Object type inheritance is very helpful to generalise e.g. similar product types in order to harmonise definitions and reduce redundancy in the model. Fig. 2 shows the basic product model of the DIMS Product Library as a UML static model diagram, using these modelling elements. Individual configurations are performed by extending the object type inheritance.

Fig. 2 also illustrates the metadata modelling process. The Product Library modelling tool allows the operator to define new modelling elements, to store them in a model repository and to re-use previously defined modelling elements. This allows e.g. to define common attributes and supertypes, and inhibits an uncontrolled growth of similar definitions in a multi-mission context. The model extension configuration process terminates with the generation of the required database scheme, automatically derived from the model repository.

Besides decoupling modelling from physical metadata storage structures, the model repository serves another very important need. Since metadata representations and models can differ across distributed systems and can change over time, models have to be described in a highly abstract manner (“meta-meta” level) to ensure adaptability. In fact, only the model of the model repository itself, the “meta-meta” model defining the structure of the modelling elements is expected to be constant. The configured product model is simply an instance in the abstract “meta-meta” model.




Fig. 2. Basic product model (extract) of the Product Library and the model repository

Tab. 2. CEOS CIP mapping of the core data model (extract)

CIP Tag Path	CIP Element Name	Path type	Core model attribute name
(4,4080)	PRODUCTDESCRIPTOR		
(4,4080)/(4,2062)	TEMPORALCOVERAGE	relative	temporalCoverage
(4,4080)/(4,2059)	SPATIALCOVERAGE	relative	spatialCoverage
(4,4080)/(4,4000)	1 { ACQUISITION } N	relative	acquisitionParameters
(4,4080)/(4,4026)	DATAORIGINATOR	relative	creation
(4,4000)	ACQUISITION	relative	acquisitionParameters
(4,4000)/(4,4001)	AQUISITIONSTATION	relative	acquisitionFacility
(4,4017)	CIRCLE	relative	boundingCircle
(4,4017)/(4,4072)	POINT	relative	center
(4,4017)/(4,4094)	RADIUSVALUE	relative	radius
(4,4026)	DATAORIGINATOR	relative	creation
(4,4026)/(4,4066)	ORIGINATOR	relative	site
(4,4026)/(4,4086)	0 { PROJECTNAME } N	absolute	mission

The metadata adaptability based on a model repository is essential for supporting interoperability and metadata standards [7]. DIMS also uses the model repository to define the mapping to the CEOS Catalogue Interoperability Protocol CIP [3]. Each applicable CIP element can be related to the corresponding attribute modelling element in the core data model (see Tab. 2). The relation cannot be generally solved using a single association, because the mapping of a CIP element depends on its context in the element path. Therefore we introduced mapping rules, also part of the model repository, to define the mapping of CIP elements and (sub-)paths. A similar mechanism is applicable to support the mapping to the ISO 19115 metadata standard [4].

The formal representation of the mapping to external standards allows building interpreters for the translation of search and present requests. As performance is of primary importance in user services, DIMS develops compilers generating direct access middleware software translating metadata requests directly into SQL requests. Thereby the compilers directly rely on the model repository and the compilation has to be repeated each time the supported protocol specification or the core data model is modified. The compiler itself is as stable as the model repository structure and the underlying protocol and description language to be served (Z39.50 and ASN.1 in the case of CIP).

Object Query Language

The Object Query Language supported by the Product Library has been developed for the multi-mission earth observation domain. It is based on the ODMG OQL standard [8] and enriched by spatial search operators and navigation along associations. The Product Library allows using OQL queries in different interfaces:

- The DIMS Operating Tool, the universal graphical human-machine interface of DIMS, for browsing product collections and entering ad-hoc OQL queries. Results are displayed in a tree-table and on the map.
- The CORBA query interface of the Product Library, used by many DIMS components to search and access product metadata in the service-to-service communication.
- The CORBA subscription interface of the Product Library, used by several other DIMS components to subscribe for product events in the Product Library. OQL is used to specify the condition that has to be fulfilled by the product so that the event is fired.
- The Product Library batch client for manual query at command line.

In all interfaces, the Product Library allows performing either synchronous queries with result sets of up to 500 entries or incremental queries allowing the client to retrieve stepwise a potentially unlimited number of hits. Query results are returned in XML representation according to the metadata model configuration. The OQL syntax follows the established pattern as defined in the ODMG standard:

```
select { projectionList } from Type [ where condition ] [ order by sortCriteriaList ]
```

The projection list specifies the attributes to be returned for the items matching the condition. If the projection list is omitted, only the identifiers of the items are returned. Aggregation functions in the projection list allow condensing the result to minimum, maximum, average and other statistical values. The item type specified in the query is the container, e.g. a product type or a product component type in which results are expected. Correlation queries over different containers are specified through implicit reference conditions explained later. The optional sort criteria list ascending/descending attributes for the order in which matching results are returned.

The condition is usually the most important and complex part of the query and unmasks the capabilities of a query language. As the full syntax description of OQL conditions fills several pages, we will concentrate on some prominent examples here. Refer to [9] for the full OQL specification. A basic condition is the comparison of attribute values and literals, and the combination in conjunctions and disjunctions. The Product Library enriches this by the following special operators to match a stored region to a spatial search area (or vice versa). (The geospatial computation is provided by the underlying IBM IDS database server [10], similar functions are provided by the corresponding extensions of most other database server vendors [11].)

- **beyond** The stored region is beyond a specified distance to the search area.
- **inside** The stored region lies completely inside the search area.
- **intersect** The stored region intersects the search area.
- **outside** The stored region lies completely outside the search area.
- **within** The stored region lies within a specified distance to the search area.

Besides these basic conditions, OQL allows searching items with a certain composition of set and list attributes. The number of values in the set/list as well as the existence of certain values and a condition on values can be used to restrict the result. The same applies for associations: The number of associated items as well as the existence of associated items fulfilling a certain (remote) condition can be specified to influence the result. The following example returns SRTM interferometric dataset products that include a certain search point and have no browse component:

```
select from SRTM1.X-SAR.DT where ( 'GeoPoint((41.2, -8.6), ANY, ANY)' inside
                                   spatioTempCoverage and number of hasBrowse = 0 )
```

A correlation of items of different types can be queried with an implicit reference condition as shown in the following example. The item type specified after the keyword ‘corresponding’ is the type to be correlated with the type in the query. The correlation condition is specified after the keyword ‘with’, and an additional condition to be fulfilled by the correlated items of the remote type can be appended after the keyword ‘where’. The example query returns the lowest and highest scene index of SRTM interferometric datasets (IFDS) on a certain data take that have no equivalent with the availability ‘existing’ and the revision ‘OPSI’. This query is used to verify consistent IFDS processing over a given data take: Normally each scene has two representations, a ‘preliminary’ and an ‘existing’ occurrence. The query returns only scene indexes for scenes having no corresponding ‘existing’ occurrence, hence where the processing is incomplete.

```
select min sceneIndex, max sceneIndex from SRTM1.X-SAR.IFDS
where ( dataTakeOrbit = 61 and dataTakeSuborbit = 10 and
       there exists no corresponding SRTM1.X-SAR.IFDS
       with equal dataTakeOrbit and equal dataTakeSuborbit and equal sceneIndex
       where (availability = 'EXISTING' and revision = 'OPSI'))
```

THE ROLE OF OQL IN THE MULTI-MISSION GROUND SEGEMENT

Processing Scenarios

Processing systems in the multi-mission infrastructure gather input products and auxiliary data and generate new output products. The way input products are identified is usually simple for lower level processing; the relationship between input and output is almost one to one and inputs can be identified by value reference, optionally with a condition and sorting to find the newest and best quality data, e.g. the best available orbit information. The input product and information determination is more complex for higher level processing, re-processing e.g. after processing algorithm enhancements, and product quality assurance.

A simplification of the processing of SRTM interferometric datasets (IFDS) is the first processing scenario example discussed. During the preparation of the SRTM mission it was decided to process the X-band mission data by continents. The result of the screening was partitioned in data takes, a data take describing all the data from coast to coast, the sensor being switched off above sea. The next step in processing was the generation of IFDS, which had to be performed consecutively through the screened data take. To accelerate the overall processing, the screening and interferometry processing was pipelined. Two queries illustrate the identification of the next IFDS to be processed (see Fig. 3): first the identification of the available data takes (DT) intersecting the continent to be processed, second the identification of the preliminary IFDS on exactly these data takes as consistent input for the next processing batch.




Fig. 3. DIMS Operating Tool visualising two queries illustrating SRTM interferometric dataset processing

Another interesting example of OQL for processing is the re-processing of ERS GOME products up to level 3 (global ozone coverage maps) in DIMS. A product management requirement states that all revisions of data have to be kept in parallel in the reference archive, since e.g. long term analysis of atmospheric processes have to be made possible on consistent data series. The selection of orbit-related level 2 input products for systematic level 3 processing has to cover the time integration interval and the new level 2 revision. In order to organise the backlog re-processing, additional query criteria divide the input data into regular consecutive portions.

To meet the very high quality standard of products in the EUMETSAT Polar System (EPS) Satellite Application Facility (SAF) network, a sophisticated quality assurance is established for the processing of Metop GOME-2 total column trace gas products in DIMS. Each product is characterised by statistical metadata about the distribution and the estimated error of trace gas density. To judge the quality of a single product, these quality parameters have to be analysed in their context within the acquisition time series. On launch, the quality assessment tool, integrated in the DIMS Operating Tool, queries all products not assessed yet and queries and plots the statistical metadata time series. After a first filtering based on relative thresholds, the operator can individually judge outlier products by viewing the individual quality parameters and the quicklook. The following query is used to retrieve the quality parameters of the products in the (temporal) area of interest.

```
select qualityParameters from METOP.GOME.TC-QA
where there exists a corresponding METOP.GOME.TC
    with equal orbitNo
    where 'GeoPoint(ANY,<startTime>,<stopTime>)' inside spatioTempCoverage
```

Dissemination Scenarios

DIMS supports systematic product dissemination in different ways. The standard procedure is the upload of product metadata and browse data to connected user service or product management systems such as DLR EOWEB, ESA MMMC and EUMETSAT UMARF. Therefore, a user service interface component is set up to organise the communication. This component is configured to upload certain product types to certain destinations. Usually only a part of the archived products shall be placed into public access. For conditional upload of products the user service interface component subscribes for product insertion events at the Product Library specifying an OQL query restricting the subscription scope. The condition in the following example ensures that only products are disseminated that have

the correct quality status and a well defined revision 'OPS2', and only if for a product this revision is not existing, the revision 'OPS1' is used instead.

```
select from SRTM1.X-SAR.IFDS
where (availability = 'EXISTING' and quality = 'APPROVED' and
      (revision = 'OPS2' or (revision = 'OPS1' and
                            there exists no corresponding SRTM1.X-SAR.IFDS
                            where revision = 'OPS2' )))
```

Ordering and Delivery Scenarios

The ordering and delivery scenarios apply various OQL queries during order processing. An important step in the order validation phase is to check whether ordered products cover crisis areas. In such cases the order has to be placed on hold, requires operator interaction and high level authorisation. To determine crisis area overlaps, the DIMS Ordering Control component places OQL queries which only return all order items outside the crisis area and which are safe for further order processing. Similar queries can be used to break down or restrict order items to different areas for which distributors and resellers have exclusive commercialisation rights.

```
select from <ProductType>
where ((uniqueId = <orderItem1> or uniqueId = <orderItem2> or ... )
      and spatioTempCoverage outside <CrisisAreaRegion> )
```

A subscription order is a special ordering method usually not available online in user service systems. Special customers such as meteorological services or mission validation partners require regular delivery of newly acquired and processed products. The Product Library subscription mechanism is also used to initiate deliveries each time a corresponding product becomes available. The associated OQL query ensures that the delivery contains only products the customer expected, e.g. covering the area of interest or from a specific processor version. As delivery of each single product may not be acceptable for the customer, the subscription mechanism can be combined with a timer mechanism. Product events are then collected during a certain period of time, before starting the delivery of these products in one package. If a consistent time series of consecutive acquisitions is of primary importance for the customer, the timer mechanism adds a time interval condition to the OQL query. Specifying the time interval lying in the past ensures that potential reception and processing delays are included and that each delivery package is complete.

For higher level products the price of an order often depends on the product content. For orders of SRTM digital elevation models the price is e.g. calculated with the help of the query 'select avg coverage from SRTM.X-SAR.DEM where spatioTempCoverage intersect <OrderArea>' and this average product coverage is multiplied with the base price and the number of scenes in the order area.

Product Management Scenarios

The domain of product management has to ensure product consistency and long term preservation. Consistency checks include the verification of complete processing. The Product Library provides a consistency analysis tool allowing to plug in mission-specific algorithms, using incremental OQL queries to iterate through the product populations and check e.g. for consecutive acquisition and processing, component existence, file sizes and metadata values. To get an overview of the population of the Product Library the product manager can generate histograms on any metadata parameter. Fig. 4 shows an example with the monthly number of products inserted into the Product Library: the histogram tool iteratively performs count queries restricting the parameter 'insertionTime' to the current month.

QUERY PERFORMANCE AND LIMITATIONS

Besides query functionality, the query execution performance is important for the acceptance of a query interface. Performance first of all depends of course on the underlying hardware and database system. Predictable queries for the localisation and identification of products, as they are issued e.g. by processing systems can easily be optimised by creating appropriate index structures in the database. For spatio-temporal searches which are very frequent in the earth observation domain, the database systems provide specialised indexing mechanisms based on different representations of the globe. For earth observation, the ellipsoid representation is superior to a flat representation with anomalies at the poles and the date line, and with variable resolution.




Fig. 4. Visualisation of the SRTM interferometric dataset population by Product Library insertion time

Optimising ad-hoc queries is more difficult and may have impact on the data model mapping. The normalised relational schema might not be always the best choice, and special mappings may become required. However the middleware allowing to abstract the object-oriented metadata model from relational storage schemes is a precondition for optimizations in the relational structure without modification of the data model. The Product Library logs long lasting queries allowing the database administrator to perform targeted index creations and other performance measures.

The most expensive queries are collection correlation queries resulting in relational join operations. On large collections such queries may become impracticable even if indexes accelerate the lookup in each collection. In such cases an explicit association of the correlated collections would allow direct linking of correlated items, instantiating the correlation query in the database. However this requires prediction of the correlation query so that the association can be included in the metadata model.

CONCLUSION

A product query interface with broad functionality is essential in a multi-mission earth observation ground segment infrastructure. Many scenarios for product processing, access, dissemination and management rely on metadata and query capabilities of varying complexity. For this purpose the DIMS Product Library allows configurable modelling and provides the Object Query Language through different interfaces, for operator interaction as well as service-to-service communication.

REFERENCES

- [1] E. Mikusch, E. Diedrich, M. Göhmann, S. Kiemle, C. Reck, R. Reißig, K. Schmidt, W. Wildegger, and W. Wolfmüller, "Data Information and Management System for the Production, Archiving and Distribution of Earth Observation Products", In *Data Systems in Aero-Space (DASIA)*, Montreal, EUROSPACE, ESA, SP-457, pp. 401-406, May 2000.
- [2] S. Kiemle, "From Digital Archive to Digital Library – a Middleware for Earth-Observation Data Management", *6th European Conference on Research and Advanced Technology for Digital Libraries (ECDL)*, Rome, Lecture Notes in Computer Science, Springer Verlag, pp. 230–237, September 2002.
- [3] Committee on Earth Observation Satellites, "Catalogue Interoperability Protocol (CIP)", Specification Release B, Issue 2.4, CEOS/WGISS, June 1998.
- [4] Internal Standards Organisation, "ISO 19115:2003 Geographic Information – Metadata", 1st edition, May 2003.
- [5] OpenGIS® Consortium, "OpenGIS® Catalog Services Specification", Version 1.1.1, OGC, December 2002.
- [6] G. Booch, I. Jacobson, and J. Rumbaugh, "Unified Modeling Language: The Summary", Rational Software Corporation, 1997.
- [7] M. Jarke et al., "Meta Modelling: A Formal Basis for Interoperability and Adaptability", in *Information Systems Interoperability*, B. Krämer, M. Papazoglou and H.-W. Schmidt, Eds., Research Studies Press Ltd., John Wiley & Sons Inc., pp. 229-263, 1998.
- [8] Object Data Management Group ODMG, "The Object Data Standard: ODMG 3.0", Morgan Kaufmann, 2000.
- [9] German Remote Sensing Data Center, "DIMS Product Library Service Manual", DLR DFD, 2004, unpublished.
- [10] IBM Informix Geodetic DataBlade Module, <http://www.ibm.com/software/data/informix/blades/geodetic> 2004.
- [11] Oracle Spatial Option, <http://otn.oracle.com/products/oracle9i/datasheets/spatial/spatial.html> 2004.

**DEVELOPING A SYSTEM TO SUPPORT CROSS-INSTRUMENT
AND CROSS-MISSION SEARCH AND RETRIEVAL OF PLANETARY DATA**

Thomas C. Stein and Edward A. Guinness

*Earth and Planetary Sciences, Washington University, St. Louis, MO, 63130
(stein@wunder.wustl.edu, guinness@wunder.wustl.edu)*

INTRODUCTION

The amount of digital planetary science data acquired by spacecraft missions has grown exponentially over the past 30 years and will continue to grow in the foreseeable future (Fig. 1). Scientists and engineers are also taking advantage of technological advances to increase the capabilities of instruments hosted by a spacecraft, leading to the generation of large and complex data sets. The growth in data volume and complexity is creating challenges in terms of providing services to the planetary sciences community for searching through such data sets and in distributing these data sets to the users. In addition, the data processing resources available to researchers have improved tremendously, which has led to an increased desire to acquire and work with these large data sets. As a result data archival and retrieval systems must be designed to support scientist’s needs to search cross-instrument and cross-mission data sets, and to provide a mechanism for electronic delivery of large data volumes in ways that are suitable for the researchers’ analysis tools.

As an example of the data volume growth from planetary missions consider the Viking Orbiter mission, in which two spacecraft operated in orbit around Mars from 1976 until 1980. The pair of vidicon cameras on each spacecraft acquired over 50,000 images or about 50 Gigabytes (GB) of data during the four year mission. These data were archived on a set of CD-ROMs of which several hundred copies were reproduced and distributed to the science




Fig. 1 The volume of the PDS Geosciences Node data holdings over the next five years is expected to increase more than 10-fold.

community. With a data set of this size, it was economical and practical for each investigator to have a local repository of the data set. In the late 1980's, the Magellan mission to Venus returned more data than all previous NASA planetary missions combined. While it was still possible to mass produce the Magellan data set on CD-ROMs for wide distribution, the cost of doing this was becoming an issue and scientists were beginning to realize that searching through and working with hundreds of CD-ROMs to find data of interest to them was a time consuming task. In the future, the Mars Reconnaissance Orbiter (MRO) spacecraft, to be launched in 2005, will generate an archive of at least 50 Terabytes (TB) during its 2 year primary mission. Clearly, it would be extremely expensive to archive the MRO data sets on physical media for distribution to each scientist, even using DVD-ROMs. In addition, few scientists would be able to invest in the resources needed to locally store and access such large amounts of data.

There is a strong need for data systems to provide the science user with the ability to search through large data sets from multiple instruments and missions and to cull out the subsets of data of interest. If users select small subsets of data, then electronic distribution and/or custom volumes on DVD-ROM, for example, become viable methods of data distribution to the requestor. As science users become accustomed to acquiring data via an online system, their requirements for services from such data systems will increase. In addition to standard search methods that focus on parameters for a given instrument, users are likely to request search capabilities that involve the science planning context of the observations, such as the objectives of single and coordinated observations, and viewing footprints of the data on a map of the planet. In this paper, we will discuss the importance of archive planning and production as a means to making the data accessible long-term to the science community. In addition, we will discuss the steps taken to produce an interactive cross-instrument support tool for data acquired by the Mars Exploration Rover Mission.

NASA DATA ARCHIVING

The Planetary Data System

Within NASA, the Planetary Data System (PDS) has the responsibility of archiving and distributing data acquired from planetary missions. The PDS is a distributed data system. Scientists and engineers working on data systems and data archiving within the PDS are associated with several institutions known as PDS Nodes. There is a Central Node that provides overall management and engineering support and individual Science Discipline Nodes for Geosciences, Atmospheres, Planetary Particle Interactions, Small Bodies (comets and asteroids), and Planetary Rings. Each of these Nodes is housed at an institution where there are a core group of scientists who are active researchers in the discipline and who also understand the nature of the data and the archiving approaches relevant to that discipline. In addition to the Discipline Nodes there are two supporting nodes, Imaging and the Navigation and Ancillary Information Facility (NAIF) that consist of engineers who are experts in these data types.

Archive Life Cycle

Online data access and distribution are part of an archiving life cycle. The illustration of an archiving life cycle presented here is based on the experience of the PDS working with NASA planetary missions. The cycle begins even before the spacecraft is launched, continues throughout missions operations, and does not end until the data sets are superseded or no longer actively used in research. The initial phase in the life cycle is when the archiving authority (e.g., PDS) and the project and its experiment teams develop a plan for generating the archive and transferring it to the archive. The archive plan describes the responsibilities of the project, experiment teams, and archiving authority. It also defines the standard science data products and ancillary data to be archived, along with estimates of the data volume and a schedule for generating the archive. It is important for the project and the archiving authority to work together early in the process in order to avoid last minute problems such as errors or omissions in key documentation, or data products that do not meet the needs of the science community.

After the archive plan is in place, the appropriate Discipline Node works with individual experiment teams to define the formats of the team's standard science data products, such as the raw data (i.e., Experiment Data Record – EDR) and the reduced data products, which are data calibrated to physical units and, if appropriate, transformed into a new coordinate system (e.g., cartographic coordinates). At this stage the two groups also design the labels or header records for the data products. The PDS has developed a standard data dictionary and archiving practices that are used in producing planetary data archives for NASA missions. This use of a common set of standards is an essential foundation for building a data system and user interface for cross-instrument and cross-mission search methods. The data

dictionary provides a consistent set of keywords for the metadata that form the bases for data searches. The archive standards help ensure that the data products are well described by the product labels and that the appropriate documentation is present. The information about the data product design and its labels is described in a formal specification document that is archived along with the data. This document is also reviewed by representatives of the user community to assess the scientific usefulness of the planned data products and to evaluate the completeness and clarity of the documentation.

Once the data products for a mission are defined and documented and the archive plan with the estimates of the data volume is complete, the archiving authority can design how the new data will be incorporated into its data access and distribution system. A series of user scenarios, which are examples of how the science community would likely search the new data, are developed by canvassing the community on its desires for accessing the data. Science experts from the Discipline Node also provide valuable input into the user scenarios. These scenarios drive the development of requirements for any modifications to the data system's database, search queries, and user interface so that the new data can be integrated with existing data resources.

Early in the mission operations and before the first data delivery, a sample data set can be used to conduct a peer review of the data set and the online data system. The peer review allows selected members of the user community to evaluate the conformance of a data set to the archiving standards and to its design specification, completeness of the documentation and ancillary data, such as calibration data, and the usefulness of any software provided with the archive.

The review will also test the user interface of the online data system. As data deliveries are received during the mission according to the schedule defined in the archive plan, they are validation and included in the online data system.




Fig. 2 The MER Analyst's Notebook navigator tool. In this tool data products are organized within a tree structure ordered by instrument, location, or time. Data products may be previewed and downloaded from this tool.

Finally, after mission operations are complete it is the responsibility of the archiving authority to maintain the data sets and the systems used to access and distribute them.

Data Access and Delivery

Smaller data sets and portions of larger data sets historically have been delivered to the end user by hard media—initially as photographic prints, then 9-track tapes and optical media. Currently a majority of PDS data are delivered electronically using Internet and Internet 2 (where available) connections, although the maximum sustained data transfer rate observed is approximately one GB per hour. A viable alternative in some cases is the “data brick” concept, in which a low-cost, large-format (200- to 400-GB) external hard drive is used as the transfer mechanism. A 200-GB data brick that requires three days for file copying and mail delivery will save more than three days over electronic distribution. In addition, some research institutions are considering imposing fees to cover the cost of increased bandwidth usage. Regardless, robust data query tools allow the user to restrict the amount of data requested and in turn reduce the amount of time required for data delivery.

Recent developments in drive technology have brought improved access to lesser-used data sets. Historically, budget constraints have precluded staging any but the highest-demand data sets on hard drive systems. Low-use data have been relegated to CD- and DVD-jukeboxes, or moved offline altogether. Higher-density drives and lower-cost technologies (such as Serial-ATA hard drives) provide a solution for maintaining larger (>100 GB) but less-well used data sets. Heavily-used metadata sets continue to be hosted on more reliable SCSI media. As demand for certain data decreases, those data are migrated to slower, less robust media servers. Tape library technologies have been considered but have been found too expensive.

A caveat in the push to larger disk formats to support the requirement for increased data storage capacity is the increased time to recover from a system fault. Experience shows that a failed RAID-5 configured 320-GB SATA drive requires approximately two hours to rebuild. In addition time required to load or restore a large RAID system is significant.

CASE STUDY: MARS EXPLORATION ROVER

Background

The Mars Exploration Rover (MER) Mission is unique in comparison to most Mars missions. Rather than a small number of instruments orbiting in a predictable manner, each of two rovers is equipped with identical science payloads that are commanded on a daily basis to perform a variety of tasks. Included in the payload is suite of more than a dozen instruments and engineering sensors and components useful for science investigations [1]. The primary mission of 90 days per rover has been exceeded and both rovers continue to acquire and return data during the extended mission.

Developing the Analyst’s Notebook

The Analyst’s Notebook concept as a user interface to data sets from rover missions was developed in 1999 to document Mars rover prototype field tests. The need for a mechanism to “replay” a field test or mission day-by-day was obvious from the start. Even during early one-week field tests scientists sometimes found it difficult to recall rationale used in making the previous day’s decisions. The Notebook captures this by collecting and organizing documentation and data from the science teams each day. In the case of MER, the Analyst’s Notebook (Fig. 2) has been updated daily during the MER Mission and is available to the MER science team for use on a non-mission critical basis. This allows the data producers to review how their data is incorporated into the Notebook.

Shortly after the MER Project was selected, a Data Archive Working Group (DAWG) was formed to address issues of data archive preparation. Members of the MER science team and the PDS worked before the start of the primary mission operations to define data formats and apply planetary data standards to expected products. The MER Archive Plan [2] was written to guide the data producers in preparing comprehensive and well-documented archives. Teleconferences are held regularly to keep all members up to date on data archive matters.

There is a fundamental difference between the MER data archive and the MER Analyst’s Notebook. The data archive must fill the requirement to provide long-term data preservation, including documentation sufficient to allow

future users to understand the data. The Notebook is an interface to the data archive and therefore contains the same information, but it also incorporates additional tools to guide the user through the data. Although the Notebook does not provide computational tools, the initial version does offer rudimentary data analysis via custom searches and browse versions of images and data plots of spectra.

A key component of the Notebook is the interface design. A set of user scenarios was developed to understand how the Notebook might be used by the science community. Prototype versions of the Notebook were created for 12 different rover field tests prior to the start of landed MER operations. Feedback—from both the science users and the Notebook production team—was incorporated into subsequent versions. The Notebook was used as the means of primary data distribution during some of the tests. Having these tests was essential to producing a usable and useful data interface.

The MER Analyst's Notebook contains science and operations data products from nine science instruments, including browse versions of the image data and plots of the spectral data. A navigation tree allows the user to browse the data in order by instrument, time, or location. (The original data products are available for download.) In addition, the mission is represented through a series of summaries organized by sol. Documentation detailing the mission, spacecraft, and instruments is accompanied by daily reports from science operations. The product overview report lists the products generated for a given sol. The rover final activity sequence plan for each sol provides an additional way to browse through the mission. An interactive search tool allows the user to perform custom queries by specifying instrument, time, location, target, filter information, and other parameters. Finally, additional resources are provided to direct the user looking for further information.

Results

The first release of the MER Analyst's Notebook is scheduled for August 3, 2004. The MER Analyst's Notebook has required approximately four person-years for design, development, and implementation. The most difficult aspect for MER has been managing the variety of data products, understanding how they are used by scientists, and developing tools to appropriately represent the data. The Notebook is a major step forward in planetary data accessibility in comparison to previous and existing tools. The presence of documentation detailing the decision-making processes is unprecedented. There are still a number of shortcomings with the Notebook however. Dynamic overlay of instrument footprints, feature and target locations, and other information is unavailable. In addition, queries are restricted to the metadata rather than supporting searches on the actual data. These types of support will be required in the future to provide a fully-functional interface to data sets too large to disseminate through historical channels.

FUTURE WORK

As multiple planetary missions return data for a given planet (e.g., over 30 years of exploration of the planet Mars), online data systems will have to continue development of capabilities for cross-mission searches, along with cross-instrument searches. In addition, there are opportunities for cooperation among international archiving agencies, for example, between the PDS and European Space Agency (ESA) archive system for the Mars Express mission. In effect, a distributed, international online data distribution system can create a virtual planetary mission where a science user can study regions on the planet using data from several missions. For the case of Mars exploration, scientists can utilize imaging data from the Viking, Mars Global Surveyor (MGS), Mars Odyssey, and possibly Mars Express; spectrometer data acquired by MGS, Mars Express, and, in the future, from MRO; and elemental composition data derived from Mars Odyssey data sets.

Technology advances have helped PDS keep pace with the expanding planetary data requirements. Multi-TB disk systems and multi-processor computers have been able to provide needed resources. Yet two factors will continue to challenge the ability of technology to provide the necessary resources: the amount of data that must be transferred to the science community and the data processing requirements levied on a data system by increasingly sophisticated (and more demanding) users.

- [1] J.A. Crisp, M. Adler, J.R. Matijevic, S.W. Squyres, R.E. Arvidson, and D.M. Kass, "Mars Exploration Rover mission", *J. Geophys. Res.*, 108(E12), 8061, doi:10.1029/2002JE002038, 2003.
- [2] R.E. Arvidson, S. Slavney, J.A. Crisp, "Mars Exploration Rover (MER) Project Archive Generation, Validation and Transfer Plan". MER 402-1-200 Rev. B, May 11, 2004.

The Planetary Science Archive

Introduction and Overview

J. Zender, G. Schwehm⁽¹⁾, C. Arviset⁽²⁾

⁽¹⁾*Research and Scientific Support Department of ESA
ESTEC*

*Keplerlaan 1, NL- 2200AG Noordwijk
Email: Joe.Zender@esa.int, Gerhard.Schwehm@esa.int*

⁽²⁾*Research and Scientific Support Department of ESA
ESAC*

*Urb. Villafranca del Castillo
P.O. Box 50727, 28080 Madrid - SPAIN
Email: Christophe.Arviset@esa.int*

1 INTRODUCTION

When the European Space Agency (ESA) approved its first long-term program HORIZON 2000 in 1983/84, it became clear that ESA needed to establish its own science data archive for planetary exploration missions. Rosetta, the first cornerstone mission took the lead in the definition of such an archive. And therefore, right from the beginning of the Rosetta mission, due consideration was given to establish such a science data archive with the view that this should be used mandatory for all future ESA planetary missions.

Currently, ESA has the following planetary missions in orbit or in the implementation phase: Huygens, Mars Express, Smart1, Rosetta, Venus Express and Bepi-Colombo. To make the best synergy of the already existing archives for the astronomical observatory missions, especially the Infrared Space Observatory (ISO) mission, it was decided to develop the archive design in close collaboration between Science Operations and Data Handling Division (SCI-SD) and the Planetary Mission Division (SCI-SB) within the Research and Scientific Support Department (RSSD) of ESA.

As most of our missions are in close collaboration with NASA, and with the experience gained in the archive design of the Giotto mission which was implemented in the frame of a large international collaboration in the International Halley Watch with data from ground-based observations and the results from the space mission, the Planetary Data System (PDS)¹ Standard was originally chosen as the underlying standard. This has the added advantage that the worldwide scientific community can work with a well-known standard. It is the aim of the Agency to provide reviewed and validated data products to the scientific community via a single interface and method independent of the instrument or mission. We introduce in this paper the architecture of the archive and describe the improvements over the current standard. Items we still believe needs further discussion are addressed at the end of the paper and provide proposals for the evolution of the standard.

2 BACKGROUND

2.1 Definition

The *Planetary Science Archive* (PSA) is the initiative, the setup, the process and the implementation to preserve data from ESA's spacecraft to planetary bodies, as well as supplementary information acquired in laboratories or ground-based observatories. The major goals of the PSA are:

- spacecraft and ground-based data preparation for long-term archiving and the preservation of these data
- distribution of scientific useful data to the world wide scientific community
- provision of data services aiming to maximize the usage of planetary mission data and ease the scientific data analysis

2.2 Overview

In contrast to observatory type (astronomy) missions, the instrumentation design and implementation as well as the data processing for planetary missions are under the responsibility of a Principal Investigator. Based on the Convention of

¹ The term 'PDS' is used for the PDS System (including setup, structure, management, finance, etc) and the term 'PDS Standard' is used when referring to the Standard Reference Document that PDS released.

the European Space Agency (1) and reflected in high-level plans and memoranda, the data from spacecraft belongs to ESA and it is the Agency's responsibility to coordinate the science operations and the data archiving. The tasks are usually coordinated within a mission Science Operations Center (SOC) implemented by the Research and Scientific Support Department (RSSD) within ESA's Directorate of the Scientific Programme.

Different to NASA's Planetary Data System (PDS), the PSA coordination is centralized and implemented within the Science Operations Centers; the data processing is done decentralized at the Principal Investigator locations similar to NASA's PDS setup.

To define the requirements of the scientific community in respect to data services, a PSA User Group has been established in which several scientific and engineering disciplines are represented: imaging, chemical analysis, radio science, radar observations, dust measurements, plasma measurements, engineering information and lander/probe coordination. The User Group specified requirements that lead to the implementation of the PSA database and data services².

2.3 Concept

In the following the concept of the PSA is discussed based on the Open Archival and Information Standard (OAIS) model (2). The term *user* was used by the PSA User Group and is synonym to the OAIS term *consumer*.

2.3.1 Consumer

The PSA User Group identified four different users. The *Standard User* – typically a scientist from the scientific community searching for reduced data ready for analysis. The *Expert User* – e.g. an instrument specialist querying for very detailed instrument or spacecraft data. The *General Public* – typically searching for the latest images reflecting scientific highlights. *External Systems* – other informational databases allowing querying and retrieval of meta-information, data or documentation.

2.3.2 Producer

Three different types of data producers are identified. The *Mission Operations Center* prepares the spacecraft housekeeping and the auxiliary data, e.g. orbit and attitude information as well as event information and command history. Any *Institute and Laboratory* might produce data sets important for long-term preservation. E.g. cometary ground-based observations from the Rosetta target Churyumov-Gerasimenko are archived in the PSA, if the data are not accessible via databases from the home observatory. Also, laboratory measurements on available samples allowing (re)calibration of spacecraft data – typically spectral databases – will be fed into the PSA. The Principal Investigator team might want to archive ground-based calibration data in the PSA that are usually not requested for ingestion.

The third type of data producer is the *Principal Investigator Team* processing all instrument data into all data processing levels, validating and ingesting the data into the PSA.

2.3.3 Coordination and Management

The coordination of the overall activities is within the Planetary Missions Division of the Research and Scientific Support Department (RSSD) of ESA. It is however clear that the coordination and management of the individual Investigator Teams is the full responsibility of the respective Principal Investigator. RSSD is however responsible for data quality and in-time delivery of the data to the consumers.

Considering the data flow from the spacecraft to the end user, the PSA archiving process implements the following functionality:

2.3.3.1 Quality Control

Quality Control is the shared responsibility of the Principal Investigator and the European Space Agency. It includes managerial as well as archiving process quality control. To ensure data quality, an external review concept process (4) was defined and will be applied to all missions.

2.3.3.2 Data Monitoring and Screening

To detect anomalies at the instrument, spacecraft or the ground-system, the data monitoring and screening function is shared by all parties: Mission Operations Center (MOC), Science Operations Center (SOC) and Principal Investigator

² Further on the term 'PSA' is used for the overall PSA system as well as for the PSA database (including Graphical User Interface, data services, external connectivity, etc)

(PI) Team. The MOC is monitoring basically the spacecraft and ground-system behaviour. Also critical instrument parameters are continuously monitored. The SOC monitors the execution of the scientific operations to ensure that planned observations finally lead to high quality data products. The PI Team is monitoring continuously the quality of their instrument and their data as well as their data processing chain. The Monitoring and Screening facilities allow the data tracking and checks for data delivery completeness.

2.3.3.3 Data Validation

The PI Team shares this task with the SOC. Whereas the PI Team concentrates on the scientific validation of the data of their particular instrument, the SOC complements this task by cross-verification between instruments and also to the science planning information. The data formats and structure of data sets are also validated by both.

2.3.3.4 Long-term Archive

The long-term archive of all mission data is the responsibility of the European Space Agency, implemented by the Research and Scientific Support Department (RSSD). In ESA we do not have an equivalent to NASA's National Science and Support Data Center (NSSDC). The way to guarantee the long-term preservation and availability of the data is to share the infrastructure with other spacecraft data. The long-term data preservation of the data contained in the PSA is therefore combined at the European Space Astronomy Center (ESAC) hosting data from e.g. the Infrared Space Observatory (ISO) spacecraft, X-Ray Space Observatory (XMM-Newton) spacecraft and data from all planetary missions.

2.3.3.5 End User Services

The distribution of the data to the scientific community and herewith the end user services are ESA responsibility. Currently implemented in the PSA are search, ordering and announcement facilities. The PSA support the concept of individual users, as well as user groups and guarantees data rights. E.g. data ingested before the end of the proprietary period are only accessible to the PI team.

3 DATA PROCESSING CHAIN

The concept of data processing level is important to understand the data processing chain. Several numbering schemes and definitions are available to define the data processing levels, e.g. the Committee on Data Management and Computation (CODMAC) definitions (3), the NASA data processing levels and the levels used by the Earth Observation Community. As none of the above-mentioned schemes define a sufficient definition for all instruments, the authors propose a more flexible scheme:

- Data Processing Level (DPL) 1a are data in telemetry format sorted and cleaned, such that they are useful for the experimenters for further processing.
- DPL-1b is equal to DPL-1a (digital number information is equal) with additional annotation and user-friendly data structures. E.g. the content of telemetry packages that make up an image are stored together in an image structure together with the necessary meta-data.
- DPL-2 is the DPL-1b data, processed to be useful for direct scientific data processing. The required amount of processing is knowledge of the science discipline and might include calibration processing, modeling processing, coordinate transformations, etc.
- DPL-3 are data products derived from or combinations of DPL-2 data products. Typical examples are digital elevation models (DEM) or image mosaics. The scientific community usually expresses the wish for DPL-3 products.

The exact descriptions for these data processing levels need to be defined for each instrument in the Archive Plan of the mission.

There are two parallel chains that can be identified in the data processing.

In the first chain, the spacecraft housekeeping parameters are processed into calibrated units. This task is done either by the MOC, the SOC or the PI team itself. As the DPL-2 data of this processing chain is needed in the L2-processing chain of the science information, it is necessary to archive these data either together with the instrument DPL-1b data or in an individual data set. As the calibration of these parameters is usually straightforward and understood even before launch, there is no necessity to archive the DPL-1b data products.

In the second chain, the instrument data are processed into DPL-1b data products. Instrument housekeeping values are normally needed for this processing step. These data products make up the basic, lowest-level long-term data products of the mission. Instrument housekeeping parameters that are needed in further processing steps are typically merged into the DPL-1b data product itself. The DPL-2 data processor is the most difficult and important processing step. Instrument calibration information, flight dynamics information (orbit, attitude, events and derived quantities) and information from scientific models are typically needed as well as a good understanding of the instrument behaviour and aging. A reprocessing of the DPL-2 data products is unavoidable in most cases. As the DPL-2 data products are the starting point for the community scientists, it is important to release these data early and in a good quality. There is obviously a great challenge and in some cases a conflict between the requirements and possibilities.

The DPL-3 data processing might include the correlation of several instrument data from one mission or even from previous missions. The complexity of the DPL-3 processing depends on the quality and the structure of the DPL-2 data. As the DPL-3 data processing is normally not defined until the end of the mission's nominal phase, a reprocessing of the DPL-2 data might be required again.

The DPL-1b processing of the scientific data is currently the responsibility of the PI team. This processing step could however be done at the SOC in the future. The DPL-2 processing is under the responsibility of the PI team to ensure the scientific quality of the data processing. Any collaborators within the scientific community – including the PI team – can take over the DPL-3 data processing.

4 THE ARCHIVING PROCESS

The archiving process follows a typical process cycle: collection of requirements, definition, design, implementation, test, active archive and the maintenance phase. The top-level archive requirements are defined in the missions' *Science Management Plan* (SMP) which is the basis for the *Announcement of Opportunity* (AO). It is then further detailed and specified in the *Archive Generation, Validation and Transfer Plan*, the Archive Plan. The Archive Plan defines also the individual responsibilities of the players, the data products and data types that are expected to be delivered. The Archive Plan shall be signed one year before launch of the spacecraft.

Each data producer shall design its data pipeline and archiving products until launch, described in the *Experimenter to Archive Interface Control Document* (EAICD).

The active archive phase follows the implementation and test phase. The active archive phase starts at the end of the proprietary period of the commissioning data. The proprietary period for ESA missions is defined to be between 6 and 12 months (5) and set by ESA in the Science Management Plan. Typically, the proprietary period is 6 months. The active archive phase ends about one year after the cease of the mission. During the active archive phase, the Principal Investigator Teams collect, process and deliver the required data to the PSA in intervals of 3 to 6 months.

During the maintenance phase updates from the Principal Investigator teams are accepted and supplementary data sets can be added to the PSA database.

The process is controlled by a review process including:

- first review phase concentrating on the Archive ICD (launch)
- second review phase concentrating on the first data delivery (commissioning + 6 months)
- third review phase reviewing overall mission archive (end of mission + 6 months)

5 THE PLANETARY DATA SYSTEM STANDARD

The Planetary Data System (PDS) Standard is a set of guidelines published in the *PDS Standard Reference Document* (3). The guidelines contain detailed information on pre-defined data types, the syntax and grammar of the Object Description Language (ODL), a data dictionary, rules on labeling data products, catalogue templates and the organization of logical or physical volumes. As a general rule, each data product must be labeled. A data product is a label linked to a data file. The label is either attached to the data file itself or detached in a separate file, a label file. Documents are classified as products and a label is required. Each self-contained unit, called a data set, must contain a default set of documentation in ASCII format, referred to as catalogue files.

The glue to all data products is the standard grammar, the Object Description Language (ODL). The ODL requires that **all** meta-data be given as a *keyword = value* pair. The keywords on the left side of the pair must be defined in the Planetary Science Data Dictionary (PSDD) (10).

The basic element to the archive is of course the data products. The standard foresees a small amount of predefined, simple, data types for these products, but allows the definition of user-defined structures or combined structures, too. Keeping in mind that the long-term preservation of the data is one of our goals, the reduced number of available data types is seen as an advantage, as it enhances the simplicity of the data on the long-term. Data producers are forced to design simple data types and data structures, which simplifies the access to the data products. To name just a few of these data types: matrix, image, spectrum, table and histogram.

The standard foresees a label for each data product. This label contains *keyword=value* pairs that describe the data type, structure and any meta-information necessary to access the data and to interpret it.

The standard also requires a set of catalogue files. A catalogue file is a label file that contains textual information in ASCII format for at least:

- a data set description (data set catalogue file)
- a spacecraft description (instrument_host catalogue file)
- an instrument description (instrument catalogue file)
- a mission description (mission catalogue file)
- all references used in the data set documentation and catalogues (reference catalogue file)
- personnel information (personnel catalogue file)
- a data set collection description (dataset collection catalogue file, optional)
- the target information (target catalogue file, optional)

Additionally, the standard requests a well-defined directory structure. All catalogue files e.g. need to be accessible from a directory named *CATALOG*. The top-level directory must contain a file named *VOLDESC.CAT* that describes the whole archive volume.

Documentation is found below the directory named *DOC* and software can be archived below the *SOFTWARE* directory. Software products might include visualisation, analysis or calibration routines.

The standard aims generally to overcome operating system dependencies and directs the archive designers towards human readability.

6 THE PSA IMPLEMENTATION

Contrary to the OAIS model describing the functional entities only within the *Archive Management*, the PSA intends to provide functionality to the *Data Producer* and the *Data Consumer*. The individual functionality that is provided or supported by the PSA is presented in this section.

6.1 Data Producer

The PSA provides consultancy to the individual data producers in form of individual meetings and teleconferences. A Data Archive Working Group (DAWG) is formed for each mission and meets regularly. In these meetings the data types, data structures and the data set organization is discussed and defined. The production of the data products is supported by several tools provided from the Planetary Data System (6). The PSA provides a toolset to validate

- the syntactical correctness of the individual data products and data sets,
- the integrity of the meta information of individual data sets
- the integrity of meta information between individual data sets
- the correct usage of units and perform value boundary checks

The toolset further collects information on the individual data deliveries and supports the data packaging and the data transmission towards the PSA. It provides data set delivery tracking and support incremental deliveries as well as updates of already delivered data sets.

6.2 Archive Coordination and Management

The functionalities of the PSA Archive itself are used by all ESA's planetary missions and are shared with the archives of all ESA's astronomy missions. The *Data Management* (Business Logic), the *Archival Storage* (*data products*) and the *Access* (User Interface) components are separated in the underlying architecture - a three-tier layer model is used. An important consideration was to separate the stored data from their final presentation to the user.

The Business Logic and the User Interface have been developed entirely in JAVA and XML, which allows the PSA to be accessible from any platform and from the most popular web browsers. This has facilitated its re-use significantly reducing development and maintenance costs (7)(8)(11).

A textual and visual presentation of the data is offered to the users to aid in selecting *data products* for direct download or retrieval through FTP. Data Sets can be retrieved through FTP. The PSA offer the provision of *browse products* or quick-look data associated to most of the *data products*. These *browse products* enable users to make informed guesses on the *data products* content and quality.

The *Ingestion* component includes the administration of data deliveries. Basically, a data delivery contains a full PDS Standard compatible *data set*. The *data producer* has the option to deliver a *data set* as sequence of *data set releases*, e.g. to deliver every 3 months an accumulating part of the observations. Repetitive information normally contained in each *data set* can be avoided using the *data set release* mechanism. This mechanism also allows the update (delete, add, modify) of already ingested *data products*.

The queriable parameters in the User Interface can contain any keyword within the *catalogue* files, the *data products* or the content of any *Index File* delivered with the *data set*. The mission’s Archive Scientist defines these parameters and the corresponding user interfaces in cooperation with the instrument teams and after feedback from the user community.

6.3 Consumer

The consumer can query the *data products* and *data sets* via the JAVA based user interface supported on nearly all available computer platforms. The query results are presented as *data products* or *data sets*, or a mixture of these. The user can immediately download individual *data products* or move selected *data products* and/or *data sets* into a Shopping Basket. User selected data are transferred from this basket into the PSA FTP area, on which the data resides for a short period of time, typically a month.

A save mechanism allows the user to save her queries and rerun them during the next visit of the PSA. The user has the possibility to directly enter the SQL query string and circumvent the graphical query panels. As a further alternative, a user can request information from the archive via an electronic protocol based on HTTP and request *data products* and/or *data sets* of interest via the same mechanism. This protocol is open and end users may wish to design their own access to the PSA.

7 INSTRUMENT/MISSION/DISCIPLINE INDEPENDENT PARAMETERS

The PSA User Interface is based on the idea of instrument, mission and scientific discipline query parameters: the user shall not be bound to a specific instrument, mission or discipline nomenclature of a parameter. The parameter definition shall be the same in whatever context it is used in the PSA. The solution to this requirement can be considered as an open problem. The underlying PDS Standard dictionary is quite inflexible and therefore stable and after a usage of more than 10 years it could be expected to be mature. The usage of keywords and its values by the Principal Investigators is however not consistent - sometimes caused by personal background and experience; sometimes consistency seems to be impossible.

To avoid potential misinterpretations in the definition of geometrical information and further ambiguities in the cross-instrument queries, an effort was done within the Mars Express Data Archive Working Group (DAWG) to define a minimal set of geometrical and positional parameters applicable to **all** instruments. Table 1 lists the parameters further categorized into top-level themes.

Table 1: Geometrical Parameters and their Classification

Category	Parameter
Position Generic	geometry epoch, orbit number
Solar Related	solar longitude, sub-solar latitude, sub-solar longitude
Spacecraft Related	spacecraft-sun distance, spacecraft-sun position vector, spacecraft-sun velocity vector, spacecraft-target velocity vector, s/c altitude, sub-s/c latitude, sub-s/c longitude,
Instrument Viewing	local true solar time, latitude of start point, longitude of start point, latitude of end point, longitude of end point, central latitude and longitude, phase angle, incidence angle, emission angle, slant distance, north pole azimuth angle, sub-s/c azimuth angle, sub-solar azimuth angle, horizontal pixel scale, vertical pixel scale

In a next step the definition of a ‘footprint’ was harmonized and agreed on so that all instruments use the same concept of an observation ‘footprint’ allowing to use one common user interface to query geometrical information from environmental sensors and e.g. mapping sensors.

8 AUXILIARY DATA SUPPORT

The data producers and potential users of the *data sets* requested additional support for ancillary data, especially the production of SPICE kernels. SPICE kernels are data files covering orbit, attitude, event, frames, mounting and alignment information (12). With support of the PDS Navigation and Ancillary Information Facility (NAIF) team, JPL, the PSA designed an automated conversion tool for orbit and attitude data for Mars Express, soon to be expanded to all other planetary missions. Other kernels were created for Mars Express in close collaboration with the individual instrument teams, the mission project team, the flight dynamics team and the NAIF team. The NAIF team offered workshops for interested instrument teams supported by the PSA.

9 AVAILABLE DATA SETS WITHIN THE PSA

In summer 2004, the PSA contains data from the Giotto 1P/Halley fly-by and ground-based observations from comet 46P/Wirtanen. The data set preparations are ongoing with the experimenter teams of Smart1 and Mars Express. The dissemination of these data will start in January 2005 in intervals of 3 and 6 months, respectively. The archive preparations with the Rosetta teams are also ongoing and the instrument commissioning data will be ingested in 2005 into the PSA as proprietary data. The Huygens archive preparations are already advanced and the Huygens data will be available for public access in mid 2006.

The PSA also supports the preservation and rescuing of already existing, non-archived data. E.g. it is planned to ingest all the Giotto data into the PSA and activities to scan image data from comet 1P/Halley from the 1980s onto digital media is ongoing.

10 OPEN ITEMS

10.1 Dictionary Related

Based on the experience gained, the authors believe that the maintenance and evolution of the PDS dictionary can be handled more efficient and in favor of the scientific data producers and data consumers. A data producer sees herself in the position to defend her request for new data structures and dictionary keywords against data engineers at the archiving authority (PDS Central Node). As negotiations are often cumbersome and drag over a long time, the data producer often prefers to use existing structures – but accepting herewith compromises towards the readability and the representation of the instrument or the representation within a state-of-the-art analysis tool. It is the authors opinion, that it is not task of the data producer to defend new ways of working or new instrumentation. A pro-active archiving authority shall flag the change and propose solutions and incorporate them into the standard well before the data producer is faced with the problem and comes up with an own solution. Several additions to the PDS dictionary in the past were based on compromises instead of good archive design.

As scientists and engineers from different scientific disciplines interpret the meaning of keywords corresponding to their scientific field, different interpretations and values for keywords in the dictionary have been used in the past and present. A typical example of a keyword causing discussions between scientists is 'TARGET_TYPE'. May-be easy for a scientist from a mapping instrument mapping a planetary body to make her selection among the proposed keyword values, it is not obvious what the target type of a neutral atom analyzer around a planetary body is. Possible selection can include 'PLANET', 'PLANETARY SYSTEM', 'PLASMA CLOUD', 'SUN' already defined in the dictionary, or new candidates as e.g. 'SOLARSYSTEM ENVIRONMENT', 'PLANETARY ENVIRONMENT', 'SPACECRAFT ENVIRONMENT', etc. The problem is *not* that one of these values might be better than another, but that with modern search engines, the *data consumer* expectations do often not match the used selection criteria. A potential solution could be to offer a *selection road map* to *data producers* and *data consumers*.

Further it is worth to mention that a data standard under direct control of one space agencies causes a usage and maintenance problem for any other space agency or national agency. The authors propose therefore NASA to consider to open the PDS Standard and its control mechanisms towards other interested space agencies. A setup similar to the Consultative Committee on Space Data Systems (CCSDS) for ancillary data could be imagined.

10.2 Related to Data Processing

In addition to the loose definitions of the *data processing levels* reported previously, the archiving authorities and the space agencies have no common policy on the final *data processing levels* of the final data sets. For a spaceborn CCD camera one can find multiple implementations of the final data products. To mention a few: Galileo is archiving raw data, calibration files, and algorithms to produce processed data. The Giotto Halley-Multi-Colour (HMC) camera archived only processed data. The camera on Mars Global Surveyor archives only raw data with limited calibration information. The camera on board of Mars Express archives radiometrically calibrated data with enough information to execute the geometrical calibration. The archive of the imager on the NEAR spacecraft contains uncalibrated images supplemented by calibration information and calibrated images. Consistency would avoid uncertainties at the *data producers* part and reduce software developments and manpower investments at the *data consumer* part.

The design and the production of the uncalibrated data products (*Level 1b*) is neither a highly sophisticated task nor in the interest of the *Principal Investigator*. Therefore it shall be considered by the space agencies, the science operation teams or the archiving authorities to take over this task and produce the *Level 1b* data itself with PI support.

10.3 Software Related

The software support policy for *data consumers* and *data producers* is unclear. It is in the interest of all involved parties to have flexible, powerful software tools in support of all activities. It is however unclear and not obvious who should develop and maintain the tools. The *data producer* e.g. might deliver software routines to read, visualize and analyse his data. After the end of the mission, it is however not in the interest nor possibilities to maintain the software for a longer period of time. It is also not in the interest and possibilities of the archiving authority to take over the maintenance of the *data producers* software. The archiving authority is finally responsible to archive the data. Examining the opposite scenario: if the archiving authority would produce and maintain a standard reading, visualizing and analyzing software, the *data producer* would be forced to adhere not only to the archive standard but also to the implementation of the software. An alternative would be to use a common data format (e.g. FITS (9)) to code the data structure and to use the archiving standard (e.g. PDS) to describe the meta-information to the data format. Software would support the data format only - and not the archive format.

10.4 Archive Process Related

The archive process is usually not part of the overall mission review cycle. This gives low priority to all archive related matters. It would be wishful that the archiving process is reviewed together with the instrumentation and spacecraft.

11 SUMMARY

ESA's Planetary Science Archive has been introduced in its concept and implementation environment in comparison to the OIAS model and in the overall framework of the science operations concept of ESA. An overview of data processing, quality control and user services was given. The archive process was described. The implementation of the PSA online system was described. The PSA can be accessed via (13). In the discussion of open items important questions were raised - influenced by either heritage or new technology challenges - that needs to be solved to guarantee proper archiving services to the scientific community. Already available data sets or the ones imminent to release were listed. The Planetary Science Archive is a new concept for European planetary scientists and it is the wish and goal of the Agency to keep the European scientific community involved in the further evolvement of the archive. The basic aim of the PSA development is to give the scientists what they are looking for in a quick and easy way.

12 ACKNOWLEDGEMENT

The Planetary Science Archive Team has too many members located at ESTEC and ESAC to mention them in the list of authors. Nevertheless each of the team members contributed uniquely to the fact that the PSA online system was released after one year of development and that we are ready now for Smart1 and Mars Express data ingestion. Therefore special acknowledgements go to J. Diaz del Rio, J.L. Dowson, D. Heather, J.-L. Hernandez Munoz, I. Ortiz, P. Osuna Alcalaya, J. Salgado, G. San Miguel, J.L. Vázquez García, A.V. Venet, K. Wirth and O. Witasse. The authors would like to thank the colleagues from PDS for their training, support and openness during the last years. Special thanks to the members of the Atmosphere Node, the Small Bodies Node and the Geoscience Node. The colleagues from the NAIF team for their brilliant support and the successful cooperation. Last not least, we would be nowhere without the trust, engagement and support of the individual members of the experimenter teams.

REFERENCES

- (1) ESA SP-1271, March 2003, ESA Conventions
- (2) CCSDS 650.0-B-1, January 2002, Reference Model for an Open Archival Information System (OAIS)
- (3) JPL-D-7669, Part 2, August 2003, Planetary Data System Standards Reference
- (4) SOP-RSSD-PR-004, September 2003, PSA Review Procedure
- (5) ESA Council, Rules Concerning Information and Data, ESA/C(89)95, Revision 1, 21 Dec 1989.
- (6) Overview of the Planetary Data System, Planetary and Space Science, Volume 44, Issue 1, January 1996, Pages 3-12, Susan K. McMahon
- (7) Optimizing the Performance of ISO and XMM-Newton Data Archives, Astronomical Data Analysis Software and Systems XII, ASP Conference Services, Vol. 295, 2003
- (8) The ISO Data Archive: Overview of Scientific Content and Use, Alberto Salama et al., ESA-SP-511, 2003, Page 3
- (9) Definition of the Flexible Image Transport System (FITS), March 29, 1999, Standard NOST 100-2.0
- (10) Planetary Science Data Dictionary Document, November 1992, JPL D-7116, Rev C
- (11) XMM-Newton Science Operations and Data Archive Page, <http://xmm.vilspa.esa.int/>, accessed 1 July 2004
- (12) SPICE Information Page, <http://naif-pds.jpl.nasa.gov>, accessed 1 July 2004
- (13) PSA Information and Access Page, <http://www.rssd.esa.int/psa>, accessed 1 July 2004

StorHouse: An Affordable Strategy for High-Volume Digital Preservation

Ann Dyball⁽¹⁾, David Clements⁽²⁾

⁽¹⁾ **FileTek UK Limited**
1 Northumberland Avenue
London WC2N 5BW
United Kingdom
Email: Ann.Dyball@filetek.com

⁽²⁾ **FileTek, Inc.**
9400 Key West Avenue
Rockville, MD USA 20850
Email: David.Clements@filetek.com

ABSTRACT

FileTek has responded to client requirements for high-volume storage management solutions since the early 1980s. As the scope and size of projects have grown from 100s of gigabytes (10 to the 10th) to 10s of petabytes (10 to the 16th), the associated requirements have become increasingly more complex. This paper describes FileTek's StorHouse® strategy and methods for accommodating the unique needs of large-scale (multi-petabyte) digital handling and preservation. In particular, it highlights and examines the following FileTek suite of “*bilities*”:

- *Scalability* – accommodation of escalating data volumes and billions of entries with no system performance degradation.
- *Survivability* – assurance that information generated by and stored with today's systems can be retrieved and maintained over future generations of storage technology advancements.
- *Adaptability and Affordability* – means to quickly incorporate and exploit higher density, lower cost technology alternatives as they emerge over time.
- *Searchability versus Accessibility* – explanation of these concepts and recommended techniques for offering optimal yet cost-efficient archive inquiry service.
- *Modularity and Flexibility* – designing for, and coping with, the range of technical options and the evolution of change with open interfaces and a ‘technology-agnostic’ approach.

The paper includes a demonstrative case study – the adoption of the StorHouse initiative by the UK Meteorological Office (the Met Office).

ABOUT OUR ACTIVITIES AT THE UK METEOROLOGICAL OFFICE

Global warming is one of the most important ecological concerns facing humanity, and the thirst to understand it has intensified. But the complexity and seriousness of global warming create special challenges for those scientists trying to understand it.

In May 1990, the Met Office began an intense investigation of the implications of global warming. The Hadley Centre for Climate Prediction and Research took the lead in this endeavour, employing some of the finest meteorological minds and computer technology in the world. Their scientists continually collect weather data from the worldwide network of automatic and human observation stations. Then by analysing and feeding the collected mass of observational data into their numerical weather prediction model, originally developed for short-term weather forecasting, scientists predict – with increasing confidence – future global warming trends.

When the challenge of managing the huge amounts of data became too demanding for their legacy database and storage system, the Met Office selected FileTek for a solution. The outcome of this project will result in one of the largest – over one petabyte in size – data management systems in the world.

Organisation Background

The Met Office, located in Exeter, England, was formed in 1854 to provide meteorological and ocean current information to mariners. By 1861 the Met Office was issuing storm warnings to ports and forecasts to the press based on observations received by telegraph from the U.K. and France. In 1922, the BBC began broadcasting the weather; and in 1936, TV stations began reporting the weather. In wartime, the Met Office is indispensable in providing accurate forecasts to forces on the ground, sea and in the skies.

In May 1990 the Met Office opened its Hadley Centre for Climate Prediction and Research. Today, harnessing the power of advanced NEC supercomputers, they are one of the most technologically advanced government organisations and a world leader in the investigation of global warming, employing about 2,300 people worldwide.

The Challenge

In 1996, six years after the global warming investigation began, the Met Office realized they had a dilemma. As the data amounts grew they faced a number of different challenges that their current technology could not handle.

According to Paul Crowley, Principal IT Consultant/Information System Strategy, the Met Office, "First, we knew that if we kept on collecting data at the same pace, we were soon going to run out of storage space. Second, the tape cartridges we were using had too low a data capacity for our needs, plus an unacceptable error rate. And third, we were using a lot of internal resources to write programs to manage the data, which was not a good use of their time."

In 1997, the Met Office decided to update their current system. They consulted with the government's Central Computer and Telecommunications Agency (CCTA) and wrote and released an Operational Requirement to potential vendors. In 1999, after a thorough and competitive process that included benchmarking and proof-of-concept presentations, the Met Office chose FileTek as the prime contractor to implement the Managed Archive Storage System (MASS).

The Solution

To manage the hundreds of terabytes of MASS data the Met Office implemented the FileTek StorHouse/RM software, which runs on a Sun Microsystems™ Ultra Enterprise™ platform using Solaris™ 2.6. For tape data storage, they implemented StorageTek™ 9840 drives and media held in a Powder Horn™ 9310 silo. The overall architecture of MASS is shown in Fig. 1.




Fig.1. MASS System Today

George Purvis, MASS Systems Manager at the Met Office said, "As we proceeded along the procurement process, we realized that the storage part of the solution was really secondary. The main issue was managing the data over time. It was crucial to the operation of the whole system."

StorHouse® permits online management, storage, and access to single rows of detail data stored on a hierarchy of diverse storage devices and media types, including tape. The ability to directly access and manage high volumes of unsummarised data without moving complete files to disk was a major factor in the Met Office decision to choose FileTek.

"Our file sizes average 80 MB," said Cowley. "To have moved multiple files to disk would have created a severe bottleneck and would have also required us to install many more hard drives. StorHouse's ability to manage vast amounts of data and get only the data we needed, not the whole file, was a major factor in choosing FileTek. This feature was key in reducing network load. Other factors for choosing FileTek and its StorHouse solution included scalability, support of existing technology, and FileTek's corporate commitment to support new technology when it becomes available. Also by having StorHouse automate our data management, we are able to free up our programming resources for other critical areas."

Cowley predicted that by 2005, the Met Office database under StorHouse control would be one petabyte (1,000,000,000,000,000s of bytes) in size. By 2010, he believes it will approach two and one-half petabytes.

"One thing we can say based on our research is that there will be global warming and climate change in the future," said Cowley. "With the technology the Met Office has harnessed, we believe that we are well on the way to understanding the magnitude of this change."

Over the last 5 years the system has grown to an average load of 400GB a day, with a variable retrieval load. The Met Office has moved its operational system from Cray to NEC supercomputers with a period of parallel running. The StorHouse system continued to load and retrieve data without a break during this transition. As technology has evolved, the Met Office has chosen to take advantage of upgrading the tape storage media from 9840 to 9940A, then 9940Bs and most recently to introduce MAID technology (Massive Array of (mostly) Idle Disk) into the storage hierarchy managed by StorHouse, such that the first months worth of data now resides on very inexpensive disk before migrating automatically to long term tape storage. At each technology change the customer has had no interruption to operations as the data moves transparently through the storage hierarchy, driven by customizable system rules.

WHAT ARE THE BIG PROBLEMS AND THE KEY REQUIREMENTS OF STORING, ACCESSING AND PRESERVING LARGE SCALE SCIENTIFIC DATA?

1. You often cannot predict how big the data will grow. Recent growth in stored digital data has been phenomenal. More digital information has been stored in the last few years than in the preceding millennia. KEY REQUIREMENT: *Scalability*
2. Assuring the longevity of digital data over time can be a real problem. Paper perishes, but has huge staying power. The BBC Doomsday project foundered because the technology used to store the data became obsolete. The choice of media was thought to be key, but the obsolescence of software is also critical. Making media resilient is not the total answer, because the software to read the media also gets out of date. This is true at application, middleware and operating system level. KEY REQUIREMENT: *Survivability*
3. Requirements change; new and leading edge products come on to the market. You cannot afford to be tied to any particular product that has proprietary formats. It is important to know that data are accessible from other software products. Technology dependence is risky and often expensive in the long term. Scientific institutions are likely to have a mixture of different technologies from different times for different purposes. Communication and exchange between these different environments is often necessary, sometimes critical. KEY REQUIREMENT: *Open Modularity*
4. Technology change is escalating: disk is getting less expensive, and new, higher density tape media is appearing. Why commit now to a technology that may well be obsolete in 2 years when something bigger, better and just plain cleverer (not to mention cheaper) is about to be launched? You need to be able to quickly incorporate and exploit higher density, lower cost technology alternatives as they emerge over time. KEY REQUIREMENT: *Adaptability*.

5. Disk is cheap...so they say. The capital cost of disk and tape reduces every year. However, it is proven that the operational cost – including manpower costs and power costs – of running huge disk arrays are much more than their acquisition cost. These operational costs have not decreased year on year. Increasing amounts of data mean that costs will continue to be an issue. KEY REQUIREMENT: *Affordability*
6. Traditional digital archiving has meant dumping data on to tape. This method of preserving could be called WORN: Write Once, Read Never. Because of this, users have been very wary of consigning data to anything other than on-line disk. Traditional data management tools have favoured disk, those technologies that have used other media have a reputation for poor accessibility. The trade off between costs, accessibility and security has always been difficult for large volumes. KEY REQUIREMENT: *Accessibility*
7. As volumes grow and digital sources become more complex, more distributed, it becomes increasingly difficult to find and extract all the threads of information spread across (or hidden within) the total Digital Resource. The need for an optimal yet cost-efficient archive inquiry service becomes more important. KEY REQUIREMENT: *Searchability*

FILETEK SOLUTION ATTRIBUTES

The FileTek StorHouse solution manages a configurable hierarchy of diverse storage devices including high performance disk, capacity disk, optical, and high performance and high capacity tape. The architecture of this “managed hierarchy” is shown in Fig. 2. Detailed system architecture is shown in Fig. 3. While the FileTek solution has numerous beneficial characteristics, the key attributes for large-scale digital preservation projects include:

Scalability – The system can manage billions of entries, as well as data volumes ranging from terabytes to petabytes and into exabytes.

Survivability – StorHouse systems support persistent archives because they can span several generations of storage technology. Stored information is insulated and maintained independent of storage devices. This means that when new technology becomes available or old technology becomes obsolete – because the data is independent to the media – StorHouse can seamlessly and transparently move the data to new media.

Affordability – Surprisingly low total cost of ownership enables implementation of large volume projects previously considered cost-prohibitive. This is because StorHouse is designed as a read-only database that can take advantage of low cost media. Where disk is used, there is a minimum of additional space needed for indexing and growth and hence reduced total disk requirements.

Adaptability – StorHouse was designed to be storage device independent. As technology advances, StorHouse is able to take advantage of the latest trends in storage devices and media. A good example is the recent introduction of MAID disk (Massive Array of (mostly) Idle Disk) into the StorHouse hierarchy. For customers like the Met Office who see cost and performance advantages in MAID, the transition and migration of data from older to newer technology is a transparent background task.

Accessibility – StorHouse/RM is a relational database with embedded hierarchical storage management. But unlike typical HSM products that need to de-stage data back to disk in order to access individual records, StorHouse retrieves only the data required no matter what media is used, and no matter how large the store. This enables high performance on new and emerging low-cost storage.

Searchability – HSM products can often be designed to access specific data. However, if you want to search an organisation’s entire data asset through billions of records, looking for a specific piece of information, the task would be difficult if not impossible. The unique design of StorHouse allows sophisticated searching tools to search the data in situ and retrieve only that data that matches the search criteria.

Open Modularity – StorHouse systems are based on open system standards using commercial off-the-shelf hardware components. Capacity is configurable so users can add just the capacity needed, when it is needed. In addition, since StorHouse is data aware, it allows the transition of the data from one hardware or software environment to another without the costly and time-consuming data migration effort traditionally required.




Fig.2. StorHouse Storage Architecture

StorHouse Configuration

Table 1 lists the major StorHouse hardware components and vendors.

Table 1. Major StorHouse Hardware Components and Vendors

Component	Description
Server	Sun Microsystems – Solaris HP – HP-UX IBM – AIX
Disk	Direct-attached RAID SAN and NAS Storage EMC Centera (including Centera Compliance Edition) STK BladeStore Copan MAID Revolution 200T
Optical	Sony/Plasmon – 5.25” Plasmon/FileNet – 12” BlueLaser DVD (2004)
Tape	STK 9840 (High Speed) * STK 9940 (High Density) * Qualstar Sony AIT/SAIT (High Density) * * WORM option available




Fig.3. StorHouse System Architecture

StorHouse Application Data Access

StorHouse manages both unstructured and structured data requirements. The system's storage manager (StorHouse/SM) supports flat file data management via an API. In addition, the StorHouse file system interface (StorHouse/RFS) provides support for applications using CIFS, NFS, and NTFS. A file may be a structure of like records (fixed or variable in format), or an unstructured singular object such as a document, an e-mail, a video, etc. Multiple objects may be grouped into a collection and stored within StorHouse as a single file with associated indexes.

StorHouse's relational manager (StorHouse/RM) uses StorHouse/SM storage management to support ANSI-compliant SQL relational database management. This supports structured data applications with non-procedural query requirements. Because StorHouse/RM also provides support for large objects (LOBs) within a row, applications may combine both structured and unstructured data into one cohesive relational framework.

Other StorHouse Features

StorHouse offers comprehensive *media management*, including automated media retirement based on age, usage, or error rates. While FileTek provides on-going device/media support that spans suppliers' roadmaps for improved densities and access speeds, it also accommodates scheduled migration to new media and complete data relocation. *Automated system-managed storage* features provide applications with a single interface to data anywhere in the storage hierarchy. These features plus *StorHouse/Control Center*, an easy-to-use Windows-based GUI tool for StorHouse system and database administration, dramatically reduce the number of manual storage administration tasks typically required by other data management systems. StorHouse *data management* includes integrated, automated backup.

With StorHouse/RM, applications have *direct SQL access to entries* (or rows) on any storage device including tape. Unlike hierarchical storage management (HSM) system technology, StorHouse has no requirement to stage files to host system disk before initiating record searches. Programmers can access StorHouse the same familiar way they access standard relational database management systems

StorHouse stores small objects as records (or rows) rather than as separate files. This enables virtually unlimited scalability and allows StorHouse to manage billions of records.

StorHouse provides shared access among multiple systems (see Fig. 3) because all data management functions are internal rather than resident in a requesting client or host system. StorHouse resources are never limited to a single application.

StorHouse/RM supports multiple data types, including scalar data, and large objects such as message text, image, audio, and video.

A continuous automated system monitor detects and logs errors or potential problems. When these problems are serious, the monitor “calls home” to FileTek service engineers, who then take corrective actions to avoid system outages and ensure high availability.

StorHouse also supports distributed processing as shown in Fig. 4. In keeping with support for large archival projects that encompass multiple locations, FileTek will support distributed application access servers and StorHouse systems. All members of a distributed scheme interconnect via TCP/IP-based network facilities.

For example, StorHouse/RFS (the StorHouse file system interface) databases and file structures may be allocated to any (or all) of the StorHouse systems. Backup copies of one system’s primary store may be maintained by other StorHouse members. As illustrated, the primary data of the leftmost StorHouse (A, B, and C) can be evenly distributed among the other systems (a copy of A is on the adjacent system, B is on the next, and C is on the rightmost StorHouse). If the primary copy of A becomes unavailable, access will be automatically routed to the secondary copy. If an entire StorHouse system is unavailable, then traffic to its entire primary store will be re-routed to the respective second copies. In this case, all of the other StorHouse systems assume a minor increase in access activity due to the secondary copy distribution scheme. No one system is overly impacted.

These features plus StorHouse’s unique suite of attributes make it the ideal digital preservation and archival storage solution.




Fig. 4. Distributed Environment

BIOGRAPHIES

Ann Dyball, Regional Area Manager for FileTek International has over 25 years in the UK IT industry primarily for UK Government agencies and major corporates. Starting with 5 years at the British Post Office, then 8 years with British Gas, 11 years with British Airways and latterly specialising the issues of very large data management with FileTek.

Ann lives in West London and is a graduate of Manchester University.

David Clements, Vice President of Solutions Marketing for FileTek, has over 25 years' experience with data and storage management systems. He came to FileTek after 12 years with Teradata Corporation, where he specialized in providing companies with hardware and software solutions for very large decision support initiatives. Prior to Teradata, David worked at Cincom Systems, one of the first independent software vendors to provide database management systems.

David lives in California and is a graduate of San Francisco State University.

ARCHIVING AND MANAGEMENT OF GEOSPATIAL DATA

Timm Ohlhof, Rainer Syffus

ESG Elektroniksystem- und Logistik-GmbH, Einsteinstraße 174, D-81675 Munich, Germany

E-mail: tim.ohlhof@esg.de, rainer.syffus@esg.de

ABSTRACT

The combination of aerial images, satellite images, digital maps, GIS data, control points, derived products such as nationwide DEM, and orthomosaic databases are the data foundation of advanced image exploitation. To handle petabytes of image information, automated production workflows and an efficient metadata information system (MIS) for imagery and other relevant geospatial data are necessary.

There are only a few operational MIS available worldwide. One of them is ESG's GeoBroker[®], which is an intelligent, high-performance solution for the archiving, management, retrieval, display, and dissemination of all kinds of imagery and geospatial information. GeoBroker[®] currently supports 53 different data types including the standardized NATO formats for imagery, elevation data, raster maps, and vector data. The metadata and vector data are stored in an Oracle database, whereas the imagery and other mass data are stored in a file system. The archive data are usually stored in a SAN (storage area network) environment on RAID systems and tape libraries. During data ingest the original data are automatically converted into standard formats like TIFF and GeoTIFF. All data in the GeoBroker[®] archive are geo-referenced, so that query results are displayed both in the form of footprints located in a reference map and in an attribute list. The geospatial data can be queried, ordered and downloaded in a LAN, an Internet or Intranet environment. GeoBroker[®] can be implemented as a stand-alone tool or as an integrated application within a GIS software package like GeoMedia or ArcView. It is in operation in several German civil, military and intelligence agencies.

INTRODUCTION

Nowadays, various Earth observation sensors and technologies are applied. In many cases, images from satellite or airborne platforms are acquired and then processed, i.e. measured, visualized, interpreted, analyzed and fused. Common to all these procedures is the generation of new, value-added information about our Earth. Some of the applications are of European or national interest, such as global monitoring, change detection, disaster management support or land mine detection. In order to take advantage of the full image information content, systematic geo-referenced archiving and retrieval, data and workflow management, standards for information interchange, image information mining (IIM) techniques, and an efficient metadata information system (MIS) for imagery and other relevant geospatial data are needed [2].

During the last few years, several types of solutions for the management of geospatial data have been developed, known as MIS or catalogue systems, geoservers, geo-shops, geospatial data portals, and geospatial data infrastructures. Some commercial products are:

- MIS: MSPIN (Delphi IMM), SMMS Spatial Metadata Management System (Intergraph)
- Archiving systems: ILS Intelligent Library System (Lockheed Martin), TerraSoar (Core Software Technology)
- Geoservers: Feature Level Data Base (FLDB)(Intergraph), ArcSDE (ESRI), GeoServer (AED Graphics), terramapserv, rasdaman, RasDB (COSIRO)

ESG's solution GeoBroker[®] has the functionalities of a combined MIS, archiving system and geoserver [1]. The paper first describes the requirements for an archiving, retrieval and dissemination system for imagery and other geospatial data. Our solution GeoBroker[®] is then presented. The system architecture is described and implementation issues are addressed. Finally some applications, e.g. the build-up and management of large GIS databases (TBytes...PBytes) and the generation of simulation terrain databases, are presented and conclusions are drawn.

REQUIREMENTS

The combination of aerial images, satellite images, digital maps, GIS data, control points, derived products such as nationwide DEM, and orthomosaic databases are the data basis of advanced image exploitation. To handle petabytes of

image information, systematic geo-referencing, archiving and retrieval, automated production workflows and an efficient MIS for imagery and other relevant geospatial data are necessary. This combined MIS, archiving and dissemination system should offer the following features:

- Data model based on international ISO/TC211 and OGC (Open GIS Consortium) standards
- Handling of centralized or distributed databases of different types, processing levels, data formats, contents, size, and availability
- User management/administration
- Multiple security level and access control (e.g. for classified data)
- Automated archiving, requesting and dissemination
- Streamlined workflow and asset management
- High-speed viewing
- Web-based delivery of geospatial information on time
- Online, nearline and offline data storage for short-term and long-term archiving
- Storage capacity for TBytes...PBytes
- Interfaces to GIS packages, ERP systems (e.g. SAP), document management systems and e-shop systems

Fig. 1 shows a typical solution for the archiving, management, retrieval and dissemination of geospatial data based on a client-server architecture with different types of clients.




Fig. 1. Typical solution for the archiving, management, retrieval and dissemination of geospatial data

GEOBROKER®

System Architecture

Fig. 2 shows the 3-tier system architecture of GeoBroker®. GeoBroker® can be used either as an application of a GIS software package or as a web application. In the first case we have utilised the API of the underlying GIS application, whereas in the second case GeoBroker® runs within a web browser connected via Internet or Intranet to a web server [1].

The metadata and vector data are stored in an Oracle database, whereas the imagery and other mass data are stored in a file system. The archive data are usually stored in a SAN (storage area network) environment on RAID systems and tape libraries. GeoBroker® can access data on a central server or on distributed servers at different locations. An interface to a raster data server, which also stores its data in a database, is under development.

The GeoBroker® service running on a GIS application server provides the full GeoBroker® functionality for an efficient data ingest and retrieval. We used an object-oriented modular software design, based on (D)COM technology. This scalable multi-user system guarantees maximum data consistency and security in order to protect the archive data from unintentional damage or intentional systems hacking. The web server software generates the HTML pages using ASP (active server pages). For the dynamic display of vector data in a web browser, an internet map server is required. To

this end, we presently use Intergraph's GeoMedia WebMap Professional. An alternative solution with an open source map server and GML (geography markup language) and/or SVG (scalable vector graphics) is planned. The web client needs Internet Explorer and the Active CGM control for the vector data display.

The system requirements of GeoBroker[®] are:

- Hardware
 - Server: PC with Pentium III 1GHz, 512 MB RAM (minimum)
 - Clients: PC with Pentium III 500 MHz, 256 MB RAM (minimum)
- Software
 - Operating system: Windows2000 or WindowsNT
 - RDBMS: Oracle Enterprise Edition 8.1.7 or 9.2, Microsoft ODBC
 - GIS viewer: *GeoBrokerViewer* (ESG), GeoMedia Professional (Intergraph), ArcView (ESRI)
 - Image Processing: ImageStation Raster Utilities (Intergraph, shareware)
 - WebMap Server: GeoMedia WebMap Professional (Intergraph)(only for *GeoBrokerWebServer*)




Fig. 2. GeoBroker[®] system architecture

System Implementation

GeoBroker[®] has an expandable open component-based client-server system architecture with an object-oriented modular software based on (D)COM technology. It runs under Windows2000 or WindowsNT with a German or English user interface and has two major software modules: the *GeoBrokerArchiver* and the *GeoBrokerWebServer*.

GeoBrokerArchiver is the core module for ingest, asset management, request, provision and dissemination of geospatial information in a company-wide network. It can be fully integrated into a commercial GIS software package, which is used for the geo-referenced display of query results and the smooth visualization of hybrid geodata in any geodetic datum and map projection. During data ingest the metadata are automatically extracted, e.g. from header data sets, and the original data are converted into standard formats (e.g. TIFF, GeoTIFF). Since GeoBroker[®] is applied for military and intelligence purposes, a multiple security level and access control has been implemented, where both unclassified and classified data can be managed. For the handling of users with different roles and rights, a sophisticated user management tool has been developed, which can be easily administrated by the customer.

Since GeoBroker[®] is a MIS, spatial or textual queries can be defined, and the query results are displayed as a hit list and graphically in a viewer showing the footprints of each individual data set geo-referenced in a world map. There is a direct link between the data set in the hit list and the respective footprint displayed in the viewer (Fig. 3).

The way of archiving the imagery and other geospatial data in a company-wide information distribution network depends on the required storage capacity, the required access times, degree of availability and scalability. A typical SAN/NAS (Storage Area Network/Network Attached Storage) infrastructure consists of one or more file and backup servers, fibre channel switches, RAID-1 disks, DVD jukebox, tape library with LTO drives (e.g. ADIC, StorageTek), backup/restore software (e.g. Legato Networker, Veritas BackupExec), possibly additional redundancy for disaster recovery.

GeoBrokerWebServer allows remote user access (Internet/Intranet) and web-based request, access, order, and download of geospatial data using standard web browsers and GeoMedia Web Map Pro for visualizing the vector data. For data provision to the end-users, the public web server is located in a non-military or public zone, because the archive and the LAN are protected against any access from external users. Fig. 3 shows screenshots of *GeoBrokerArchiver* and *GeoBrokerWebServer*.




Fig. 3. *GeoBrokerArchiver* (left) and *GeoBrokerWebServer* (right)

The following data formats are supported by GeoBroker®:

- Aerial image data: TIFF, GeoTIFF
- Satellite image data: Ikonos, Quickbird, EROS, SPOT, Landsat, IRS, ERS, JERS, Radarsat
- Raster data: ADRG, CADRG, ASRP, CRP, USRP, RLE, MRG, KMRG, GeoTIFF, Mil-GeoTIFF
- Vector data: VPF (VMap, FFD, DNC), S57, DFAD, DLM, DGN
- Elevation data: DTED, GTOPO, DHM, DSM, SRTM
- Simulation data: OpenFlight, SIF, SEDRIS (in preparation)
- Point data: TP (trigonometric point), NavP (navigation point), MP (military point), GCP (ground control point), RP (reference point)
- Maps and charts: atlas, topographic map, aeronautical chart, marine chart, thematic map, city map, world map, continent map, regional map
- Other data: slide, video sequence (MPEG, AVI), dossier

New data types and formats can be easily integrated using the GeoBroker® wizard.

APPLICATIONS

During the past few years, ESG has customized GeoBroker® for several civil, military and intelligence agencies. Typical applications of GeoBroker® include imagery archiving, retrieval and dissemination, the set-up and management of large geospatial databases (TBytes...PBytes), the generation of simulation terrain databases and decision support during civil and military crisis operations. GeoBroker® is the central catalogue system of the Bundeswehr Geoinformation Office for all imagery and geospatial data. In the following, two examples are described in more detail.

Geospatial Intelligence

An efficient data management system is a fundamental part of all geospatial intelligence solutions, where military, intelligence and homeland security organisations combine data from different sources into a single environment for better decision-making, crisis management, situational awareness and surveillance. Fig. 4 gives an overview of the many sources, services and products for geospatial intelligence tasks. Within the services, main components are:

- Generation and update of a cross-border or even worldwide geospatial database (vector and raster)
- Generation and update of target and simulation databases

- Exploitation of commercial and classified imagery
- Terrain analysis and display
- Rapid mapping, production on request
- Generation of a common relevant operating picture (CROP)
- Generation of a unified environmental description or a consistent environmental picture (CEP)
- Automated data archiving, management and dissemination
- Data sharing across the enterprise and with other agencies and countries
- Quality control of all production workflows




Fig. 4. Geospatial intelligence sources, services and products

A consistent, up-to-date, cross-border or even worldwide geospatial database is the core of all geospatial intelligence solutions and a prerequisite for the military and civil forces conducting joint, interagency, and coalition operations. The need is to collect, manage, integrate and fuse data from a large number of different sources, which have different formats, projections, resolution, accuracy and reliability. Fig. 5 shows the geospatial intelligence workflow including GeoBroker® as the central data management tool.




Fig. 5. Geospatial intelligence workflow including GeoBroker®

Database Generation and Mission Planning

A geospatial database forms the basis of any simulation system or mission planning system used in the armed forces. Examples are training simulators for flight and tactical operations of military aircrafts and helicopters, operations research (OR) simulators, mission rehearsal, and mission planning systems for unmanned vehicles (UAV) or cruise missiles.

Our central mission planning system consists of the following components:

- 5 seats with workstations and a central archive system
- Collection, management and provision of geospatial data using GeoBroker®
- Workflow-based generation of the mission databases
- Workflows for vector and elevation data processing using GMPro components
- Generation of 3D models for image based navigation (IBN) and target tracking

Due to high requirements for security, availability and reliability, the archive system for the storage of the geospatial and mission databases is redundant and installed at different locations within the military area. The archive system consists of a server cluster (RAID-1) connected by a 1 Gbit fibre-channel network to redundant SANs (RAID-5), switches, and a common LTO backup system.




Fig. 6. Central mission planning system – hardware architecture with workstations and archive system (left), database generation software with improvement of vector data and generation of 3D models (right)

CONCLUSIONS AND OUTLOOK

The efficient management of geospatial data with fast access to the data archive and the delivery of the right information at the right moment to the proper user is an essential task for many time-critical military and civilian applications. GeoBroker®, as a combined MIS, archiving and dissemination system for imagery and other geospatial information, fulfills most of the user's requirements and has been in practical use in military, intelligence and civil agencies for several years. GeoBroker® can be easily customized and tailored to user needs, existing hardware investments, and software environments. It is available either with a German or an English user interface. Due to its ease of use, GeoBroker® can be handled even by employees without particular GIS or IT expertise.

In future, it is planned to add a knowledge database to GeoBroker® as well as to appoint rules, to design models and methods and to deploy software tools for an IIM workflow, which assist the operator in feature extraction, target recognition, change detection or other image analysis and IMINT (image intelligence) tasks. Other future enhancements comprise:

- Improved viewing of vector and raster mass data
- Interface to a raster data server in order to store the raster data in a DBMS (e.g. Oracle) instead of a file system
- Interfaces to e-Shop and ERP systems (e.g. SAP)
- Integration of GeoBrokerWebServer into geo-portal and geo-infrastructure solutions based on web technology

ACKNOWLEDGEMENTS

We would like to express our thanks to our team colleagues Elke Weisgerber, Franz-Peter Gschwendtner, Bodo Probst, Reinhard Pohl, Christian Stätter, Dieter Hennecke and Albert Bauer, who provide valuable support to our projects. The proof-reading of Jürgen Weis and Alexandra Carter is gratefully acknowledged.

Parts of the work related to GeoBroker[®] were funded by the German Federal Office of Defence Technology and Procurement (BWB).

REFERENCES

- [1] GeoBroker[®], product information, <http://www.esg.de/en/services/products/geobroker.html>.
- [2] W. Mayr and T. Ohlhof T., “Geoinformation via parallel image engineering and image information mining”, *Theory and applications of knowledge driven image information mining, with focus on earth observation*, ESA Special Publication no. 553, 6 pages, Madrid, March 17-18, 2004.

“The High Speed ESA Earth Observation Network Project For Improving Data Exchange and Access to the User”

PV-2004, 5-7 October 2004, ESA/ESRIN, Frascati, Italy

Massimiliano CICCARELLI

DATAMAT S.p.A. - Via Laurentina, 760 - 00143 Rome – Italy

Phone: +39 06 5027 2515 (Direct) +39 06 5027 4570 (Secretary) - Fax: +39 06 5027 4500

E-mail: massimiliano.ciccarelli@datamat.it

Maria Eugenia FORCADA ARREGUI

ESA/ESRIN Via Galileo Galilei snc – 00044 – Frascati – Rome – Italy

Phone: +39 06 941807657

E-mail: Eugenia.Forcata@esa.int

Abstract

This paper resumes performed activities as well as future benefits and challenges in deploying a new high-speed network infrastructure connecting the major European Space Agency (ESA) e non-ESA satellite operational sites to support huge Earth Observation (EO) satellite data transfer. The activity was carried out in the frame of the **High Speed ESA Earth Observation Network (HiSEEN)** project developed by DATAMAT S.p.A. on behalf of the European Space Agency in the frame of a more general Earth Observation project named as OXIGEN (Open and Operational Ground Segment) proposed by the ESA Earth Observation Directorate. Proposed high-speed network infrastructure fully relies on GÉANT network, which is the high-speed trans-European backbone used by European research institutes. This project will be complemented soon by second step that foresees installation of high-speed firewall within the HiSEEN infrastructure.

1. Introduction

Because EO missions has demonstrated in the last years their utility in different areas of applications, ESA needed a more effective system to satisfy the users expectations in easy and quick access to Earth Observation products in a suitable format that matches their needs. As result, the Earth Observation Ground Segment Department identified following drivers in OXIGEN (O₂) project development. These drivers also naturally applied to the HiSEEN project being this latter focused on the basic network infrastructure realization.

- *Provide data quicker to the users, reducing the waiting time from ordering until data is available at user premises;*
- *Increase and facilitate access to the EO products;*
- *Spin-off new marketing distribution technologies and strategies, developed in other sectors and now strongly consolidated, in the satellite EO ground segment;*
- *Maintain or reduce operational costs in the facilities on-charge of products distribution due to internal data shipment between the different centers of the Ground Segment;*

According to above specifications, the goal was to create a network of **On-line Archive Centres (OAC)**, able to archive and electronically distribute Earth Observation data to end-users by inter-connecting different On-line Archive Centres all around the Europe by means of an **On-line Archive Network (OAN)**.

The **High-Speed Intranet Network (HSIN)**, provided as main deliverable in the frame of the HiSEEN project, will be one of the basic architectural elements composing the OAN fully relying on the GÉANT network.




Figure 1 - Summary of GÉANT Infrastructure

Hereinafter is the complete list of all ESA and non-ESA operational sites that have been incrementally connected to the new High Speed Intranet Network to finally compose the whole set of European EO sites as interconnected throughout the European GÉANT backbone as summarised by Figure 2.

- Farnborough (UK)
- Oberpfaffenhofen DLR (DE)
- Neustrelitz DLR (DE)
- Matera (IT)
- Maspalomas (ES)
- Toulouse (FR)
- Svalbard, via Tromsø (NO)
- Kiruna ESRANGE (SE)
- Kiruna Salmijarvi (SE)
- ESA/ESRIN (IT)




Figura 2 - EO Sites GÉANT Backbone Inter-Connection

2. High-Speed Intranet Network: Specifications

This section copes with most important specifications to be satisfied by the new high-speed network infrastructure. It shall allow interconnections among main European Earth Observation centres permitting electronic data shipment between ENVISAT, ERS-1/2, Third Party Missions and future locally managed satellite missions (e.g. GOCE, Cryosat, ALOS, etc) as well as a remote Data Distribution System (DDS) usage. By the same time it shall also allow the virtual access to the OAC through the User Data Access Portal (User DAP) as well as the data distribution between the acquisition and processing and archiving centres, the movement of data from the Data Archives to the best Distribution Node to the IP/DVB-S Transmitting Station for allowing distribution over satellite from Distribution nodes where no uplink is available.

The new high-speed network infrastructure shall also allow an user to access, according to its own user profile, to different FTP servers interconnected to the INTERNET by means of different INTERNET access points. Such accesses will differ each other by their own characteristics because relying on different commercial and/or academic links (i.e. actually GÉANT). The main constrain in performing what above stated was to ensure that, once accessed to an host via a dedicated and specific network path, both forward and reverse paths shall rely on the same routing tree so preventing any not symmetric routing between the end-user and the selected host.

Finally also an investigation about actual Multicast applications suitable for supporting reliable satellite data distribution around European EO sites was performed: current Ground Segment data distribution policy foreseen in fact different Off-Line European centres to receive same data from one or more source ESA stations. The potential advantage could result in a more efficient usage of the available network bandwidth if demonstrating data dissemination among ESA EO sites, by means multicast data transfers, may be competitive respect to the uni-cast.

The set-up of the new network communication infrastructure, which shall be deployed in parallel with the current and local network infrastructures already put in place at each above ESA EO site for avoiding loose of operational services, shall be achieved by connecting the above centres to the GÉANT backbone according to local availability and specific centre requirements or constrains, if any. These new network connections shall rely on the following data security policies, supporting IP network in a native way for any existing application or any new application to be developed for e-data shipment.

- Data encryption, to avoid unauthorised usage of the EO products transferred through the GÉANT backbone;
- Protection from malicious intrusion and damage.

Last, but not least, the proposed network infrastructure will also support, on a site-by-site basis, internal LANs harmonization between locally and different managed satellite data missions. The scope is to integrate and optimize different LANs initially and historically devoted to different EO satellite data missions management into a unique site internal LAN accessible via the new established GÉANT backbone through out a re-engineered on site front-end LAN.

The following drawing summarises different OAN architectural elements, including the High-Speed Intranet Network role in supporting EO electronic data exchange between On-Line Data Centres relying on the GÉANT backbone, so the public INTERNET. Remote data providers, (i.e. entities that are not part of the ESA Ground Segments, but will only provide a specific set of EO data to be distributed through this European on-line distribution network) shall use the network to integrate data into OAN. This network will also be used to allow quick data distribution between acquisition and processing centres and to move the data from any Data Archive to the satellite transmitting station, for its distribution over satellite link.




Figure 3 - Physical On-Line Archive Network Architecture

3. High-Speed Intranet Network: Proposed Architectural Solution

Proposed architecture solution was composed by the following network infrastructure elements and SW elements which have been deployed at each site for integrating with existing LANs and for connecting to the GÉANT backbone. The following drawing summarises the generic network architectural solution installed at each involved EO site. Specific architectural solution elements as well as deviations, if any, have been derived and detailed on a site-by-site basis according to preliminary performed sites surveys results.

- a CISCO router as gateway to the GÉANT backbone and for inter-connecting local existing LANs on a site by site basis.
- a NAS host used as storage & forward mechanism for performing satellite data missions electronic shipment;
- SW Monitoring Tools gathered from the public domain world (i.e. MRGT, NTOP, NAGIOS, etc.) aimed to perform both network and electronic data shipment monitoring while configured on a site-by-site basis.




Figure 4 - Proposed Architectural Solution

As indicated by the above drawing, per each surveyed EO site usually more than one INTERNET access was available: typically at least an already existing commercial link plus the new GÉANT academic link to be put in place. As result, a Multi-path User Access feature has been locally implemented at each EO site using hereinafter fully configurable mechanisms defined at network equipment level, without requiring any re-engineering of existing on site LAN layouts for avoiding any impact on operative scenarios.




Figure 5 - Proposed Solution For Multi-Path User Access

- *Static Routing Policies*: to define default routing policies;
- *Network Address Translation Capabilities*: to translate private in public IP addresses;
- *Policy Based Routing via Route Map commands*: it allows network traffic to follow a particular path through the network despite of the local defined default destination rules (i.e. static routing) by setting the next hop on a source IP address basis criteria;
- *Two Private & Public IP Addresses Used For Each Host*: to allow user to select between different INTERNET ISP and for being able to select return path on the basis of different private IP.

To cope with security requirements the CISCO router has been equipped with IOS supporting the following features and interconnected to the local NREN according to the proposed table.

- FW/IDS – Such feature implements a state full Firewall with an embedded Intrusion Detection System aiming to prevent external attacks. It can be used to manage authentication on incoming connections;
- PSEC/3DES – IPSec supports the Triple DES encryption algorithm (168-bit) in addition to 56-bit encryption. Triple DES (3DES) is a strong form of encryption that allows sensitive information to be transmitted over un-trusted networks.

Country	Site	Access Router	Interface
Italy	Frascati	Cisco 7304	ATM
Italy	Matera	Cisco 3745	FE
Sweden	Kiruna Salmijarvi	Cisco 3745	FE
Sweden	Kiruna Esrange	Cisco 3745	GE
Germany	Oberpfaffenhofen	Cisco 3745	E3/G.703
Germany	Neustrelitz	Cisco 3745	E3/G.703
UK	Farnborough	Cisco 3745	FE
Norway	Svalbard	Cisco 3745	E3/G.703
Norway	Tromsøe	Cisco 7206	GE
Spain	Maspalomas	Cisco 3745	FE
France	Toulouse	Cisco 3745	FE

Figure 6 - NREN Connection Interfaces

About Multicast Analysis Results, as general comment, the main issue was related to the detected dependency of the overall data transfer performances from the lowest throughput link capacity.

The more the network affected by the multicast data traffic is balanced, with respect to the available bandwidth capacity, the better is the resulting overall data traffic performance and as result the advantage of using the multicast for data transfer. Although the multicast transfer it might be of some benefits to perform the data dissemination among ESA EO sites, some issues, most of them not linked to the measured performances or reliability features, may be faced or it shall be dealt with before carrying out the multicast implementation into the HISEEN infrastructure.

Deeply analysis on required effort to on the following points shall be investigated before moving to multicast implementation over existing and already operating EO G/S:

- Configure existing network devices of each ESA EO site involved to be able to manage the multicast data transfer;
- Integrate the multicast applications with at least the ENVISAT PDS;
- Manage difficulties in the service maintenance from the personnel;
- To configure, tune and operate proper tools for traffic data management for getting a minimal assured bandwidth.

4. High-Speed Intranet Network: Integration, Verification & Validation Approach

According to identified Test Items, only one Verification Phase was identified.

- Phase 1. **“Generic Site” Network Infrastructure Verification & Validation**, whose main focus is directed onto the delivered equipment: hardware and software installed on it. This Verification Phase is repeated for each site. Validation aims to demonstrate that multi-mission data e-shipment service works properly integrated with the HiSEEN network infrastructure.

The phase was then organised per Steps, starting from On Factory then terminating with On Site activities.

- Step 1. **On-Factory Deployment and Integration Test**, encompassing the verification activities focused onto the correctness of the procured HW and SW as well as onto the right functioning before the delivery at the final destination site; this in order to anticipate problems or anomalies, if any. It is performed at DATAMAT S.p.A. premises; its output is the On-Factory Deployment & Integration Test Report.

Step 2. **On-Site Deployment Integration, Verification and Validation Test**, aimed to check the connectivity, functionality and performances at the final site. Specific tests are performed with other sites already connected to GEANT backbone before moving to the End-To-End testing session. This step is performed at each target site; its output is the On-Site Integration, Verification and Validation Test Report. Both the performance requirements' sets, specifying respectively the Transfer Speed (Mbps) and the Data load (GB/day) as well as the security requirements are in the scope of this test step;

Step 3. **End-to-End Acceptance**: whose main goal is the verification of the EO data e-shipment. The activity is used to get performance figures according to the identified performance requirements. It is executed at ESRIN premises; output of this activity is the End-to-End Acceptance Test Report. Both the performance requirements' sets, dealing respectively with the Transfer Rate (Mbps) and the Data load (GB/day) are in the scope of this Step

The selected utilities/tools used for executing functionality/performance tests were the following:

- **Ping , trace-route , ftp and telnet** – These utilities were used for test connectivity and security policies testing purposes;
- **pchar** - The tool was used for testing the average bandwidth on each hop of the HiSEEN infrastructure and the GEANT network in which pass data sent during HiSEEN connection. With this tool was possible to verify the correctness of the HiSEEN implementation respect to the expectation. It is a tool to characterize the bandwidth, latency, and loss of links along an end-to-end path through the Internet. It is based on the algorithms of the "pathchar" utility written by Van Jacobson, formerly of Lawrence Berkeley Laboratories;
- **iPERF** - The tool was used for testing the maximum bandwidth on the end-to-end HiSEEN connection. It allows to measure maximum TCP bandwidth, allowing the tuning of various parameters and UDP characteristics. The "iperf" reports bandwidth, delay jitter, datagram loss. It is also a performance measurement tool for TCP/IP and UDP/IP that provides network throughput information and statistics on jitter, packet loss, and maximum segment and transmission unit sizes.

5. Project Activities & Status - Implementation & Deployment Plan

The project main activities will be performed according to the following high-level plan:

- EO Sites Surveys;
- Reception of material at DATAMAT S.p.A. premises;
- Configuration, set-up & preliminary equipment tests;
- GEANT Connectivity via local NREN (Only for Kiruna) & OnSite Installation & set-up according to ESA agreed priority policy;
- On Site Existing LANs Integration (e.g. ENVISAT LAN, if any, ERS/TPM LANs);

Concerning network connectivity to the GEANT backbone, at the time of writing this document the HiSEEN project is assumed to be fully terminated by the end of 2004. Following tables resume, at the time of writing this document, the overall HiSEEN project status.




Figure 7 - HiSEEN AIV Approach

Country	Site	NREN	NREN Service
Italy	Frascati	GARR	OK
	Matera		September
Sweden	Kiruna Salmijarvi	SUNET	OK
	Kiruna Esrange		OK (via Salmijarvi)
Germany	Oberpfaffenhofen	G-WIN	OK
	Neustrelitz		October
UK	Farnborough	UKERNA	OK
Norway	Svalbard	UNINETT	OK (via Tromsø)
	Tromsø		OK
Spain	Maspalomas	RedIRIS	October
France	Toulouse	Renater	September

Figure 8 - NREN Connection Status

Country	Site	NREN	Site Operational
Italy	Frascati	GARR	June
	Matera		October
Sweden	Kiruna Salmijarvi	SUNET	June
	Kiruna Esrange		June
Germany	Oberpfaffenhofen	DFN	June
	Neustrelitz		October
UK	Farnborough	UKERNA	June
Norway	Svalbard	UNINETT	June
	Tromsø		June
Spain	Maspalomas	RedIRIS	October
France	Toulouse	Renater	October

Figure 9 - EO Site Pre-Operations Plan

6. Conclusions

Despite of the European Earth Observation industry annual revenue showed a growth of few percentage points in 90's, one of the most interesting aspects which should focus our attention is that, while the private market – mainly focused on the telecommunication sector - has been continuously strengthening it remains much smaller than the public sector market, mostly focused on defence. There may be several reasons for this disappointing performance:

- Space agencies worldwide put their focus on the development of space HW rather than services;
- The barrier of data cost;
- The hurdle of data access.

The HiSEEN project – and much more the OXIGEN (O2) project as whole – may represent a new perspective for Earth Observation by providing a new high-speed network infrastructure for facilitating and reducing data costs and make easier and easier the users access to multi-missions European Earth Observation On-Line/Off-Line archives. At the same time it may provide the necessary infrastructure for the development of new services which – assuming the existent communication links – could not been realized because lack in performances.

Integration, Verification and Validation tests executed between already connected EO sites fully confirmed – even if the solution still relies on an academic communication infrastructure - the practical possibility to move definitively the actual ESA Earth Observation data shipment policy from media to an electronic medium in the short term.

At the same time, it also highlighted the advantage coming from putting in place a modular and expansible network infrastructure which may allow in the short term users to remotely access ESA multi-missions long-term archives in a fully transparent and efficient way.

7. References

- [1] www.cisco.com
- [2] Developing IP Multicast Networks Volume 1 – Beau Williamson, ISBN: 1-57870-077-9 www.ciscopress.com
- [3] Enhanced IP Services for Cisco Networks, Donald C. Lee, ISBN: 1-57870-106-6 www.ciscopress.com
- [4] <http://www.psc.edu/networking/projects/tcptune/> - Enabling High Performance Data Transfers – Author: Matt Mathis <mathis@psc.edu> and Raghu Reddy <rreddy@psc.edu> (with help from *many* others, especially Jamshid Mahdavi)
- [5] www.caida.org - Cooperative associations for Internet Data Analysis
- [6] www.ncne.org - National Center for Network Engineering
- [7] <ftp://ftp.pbone.net/vol6/ftp.redhat.com/pub/redhat/linux/6.2/en/doc/HOWTOS/other-formats/html/Multicast-HOWTO-html/Multicast-HOWTO.html#toc2> - Multicast Over TCP/IP HOWTO
- [8] www.web100.org - Provide the software and tools necessary for end-hosts to automatically and transparently achieve high bandwidth data rates (100 Mbps) over the high performance research networks.
- [9] <http://people.ee.ethz.ch/~oetiker/webtools/mrtg/> - MRTG – Multi Router Traffic Grapher
- [10] <http://people.ee.ethz.ch/~oetiker/webtools/rrdtool/> - RRD – Round Robin Database Tool
- [11] <http://www.steveshipway.org/software/> - ROUTERS CGI for MRTG

“Virtual Access to Information: an emerging concept”

5-7 October 2004, ESA/ESRIN, Frascati, Italy

Jean Pierre Antikidis

CNES

“Space Information System”

Direction for Strategy and programmes

18 Av.E.Belin 31401 Toulouse Cedex 9 (France)

Email:jean-pierre.antikidis@cnes.fr

1) INTRODUCTION

The technological evolution of major human activities has been continuously oriented toward higher and higher level of "abstraction". As an example Blue-prints for creation of planes or building is fully replaced by virtual conceptual design, electronic circuitry mock-ups is now first designed with the help of SPICE simulators, point to point computer access have been replaced by Internet virtual connectivity.

A new evolution is now underway that pave the road to promising new concepts called "Virtual Information Management" in which computers systems will be replaced by virtual equivalents with no need to know anymore where the information stay and what computer produced it. The advent of GRID technology is a first step in this direction but as for Internet and the IP protocol it is only the first stone of a more larger (but ought to be much easier to manage) environment combining to-day present telecommunications progress as well as advanced computer and much higher level operating systems (some call them "Upperware").

Internet is a precursor of such a new technical and societal change that could lead soon to a full "dematerialization" of the treatment of the information enabling even small entities (commercial or privates) to create even large scale services by manipulating "virtual definitions of their needed systems and archives"

Methodology trends

From Real =>	to « Virtual »
Blueprint plans (1850)	CATIA (1985) (virtual mockup)
Formal programming (1948)	WINDOWS (1990)
MINITEL Physical links(1970)	INTERNET (1990)
Electronic circuitry (1914)	SPICE simulators(1985)
Processing systems (1960)	GRID+VIRTUAL “Upperware” (2002)

Fig.1 Computing trends

2) VIRTUAL STORAGES AND THEIR RELATION WITH GRIDS SYSTEMS

2.1 Is GRID an end or a start ?

The advent of GRID is opening the door to completely new concept in term of computerized system implementation. As usual when a new technological step forward occurs, ideas are first concentrating on the straightforward use of the new tool. This is the to-day situation with GRID mostly considered as way to produce flexible and powerful computer resources in a kind of decentralized fashion.

Therefore after a while new advanced ideas are emerging in which GRID is to be considered as a way to pool in an invisible and easy way heterogeneous systems including their associated data stores. This can be considered as the next step of a conceptual evolution in which computerized systems will become fully virtual. It could be achieved theoretically without any GRID but in practice GRID is the most elegant and efficient way to provide the necessary power.

Let's make an analogy: to make a car one need a motor. This could be a steam engine or a Gas engine. Without a motor you have no car, but a car is not an evolution of a motor. Same stands for GRID technologies: To create the next generation of "virtualized processing systems" you need a processing force. The GRID is one of the possible motor and likely to be the dominating one. But it is of high interest to analyze this new approach in the frame of a new "vision» as was the car or the plane concept with respect to gas motors developments.

2.2 The GRID Cookbook

Most of the specialists have understood GRID is a very powerful and effective tool, but also need to be complemented by some more support as seen from the user side. Up to now GRID is seen from the developer side as a set of primitives called "GRID middleware primitives". The standards are continuously evolving, but the solutions proposed for the use of GRIDs can be grouped into several classes of solutions:

A. Direct application plug into a GRID

This case has been the first being implemented and is still applicable to specific cases where the resources needed are of primary concern. This justifies the development of specific interface software ("Application tailored"). For obvious reasons it becomes expensive and inflexible for unspecialized users or quick developments needs.

B. GRID interface building blocks & GRID Thematic interfaces

This case is a normal evolution by providing a higher level of services. The middleware primitives are used in the frame of an integrated interface making use of Web based services. This allows several advantages:

- The user sees the system through a Web interface
- Submission of jobs to the GRID is also performed with the help of a Web based interface

Many of the on-going research programs are making extensive use of these features. They may be offered as "building block elements" API like calls or Thematic services layers. Therefore the complexity of the GRID handling is still present as well as the need to carefully map the data flow onto the elementary pieces of software services. One may still consider it is an intermediate evolution (like DOS with respect to WINDOWS or Object programming in the PC world) since more advanced (but less proved!) concepts are beginning to emerge.




Fig 2. "Virtual archive topology" out of a GRID system

C. GRID + Virtual High Level Operating System compounds (Upperware)

This is an emerging concept; it has been referred to in several prospective presentations concerning GRID evolution and aiming at “making GRID easy to use for most of the European users.

The idea is to make use of higher level descriptors and leave a specific operating system to take care of all necessary actions connected to the handling of GRID middleware primitives. Such an upper layer could be called “Upperware” and bring a novelty in the GRID objectives since not aiming any more at final user satisfaction but as a way to simplify and “hide” the GRID properties to the developers.

This could be compared to a high level Operating system in which the simplicity for developing applications is drastically enhancing the efficiency of the systems developers. At that stage the GRID specificity is not seen anymore. (Looks comparable to PC/windows O/S where you displace icons instead of describing the exact location of a data set as required in a DOS environment)

3) THE ADVENT OF WAGs IN ASSOCIATION WITH VIRTUAL UPPERWARES




Fig3 Wide Area GRID (large scale integration of CPU, Storage and Communication needs) being used as a “Virtual machine”

In fact a GRID system is basically a set of decentralised systems being “organised “ with the help of a suitable Operating system taking care of any materiel induced constraints. The natural evolution of such “Middleware” will encompass not only “processing activities” but also long time storage of the information. Such improved and much capable access will free us from the need of taking care of the practical handling of the stored information. This could allow unprecedented new possibilities in term of flexibility and easy creation of data bases profitable to: research centres, added values companies, support to developing countries or emerging European economies. It should also

trigger interest in and a much wider use of space-borne data by small industrial entities able to create and rapidly test access to data bases (at much lower costs) , thus inducing profitable services.

Theses approaches may also induce the need for General purpose Enabling Service able to create processing and archiving systems treating Earth Observation data and conventional data without any deployment of dedicated Hardware hosting structure. ("Virtual application platforms concept"). It may constitute a service "superstructure", plugged onto existing GRID engines, allowing any user to create its "**Virtual dedicated operational service**" with the help of high-level descriptors.

To day there is most of the time a direct relation in between data base location and data base content and when distributed data bases exist they are most of the time interconnected in a specific way. The new GRID concept will allow the use of diffused storage systems with most of the time a large and fully distributed redundancy.

These particular features are even more promising in case one make use of a large scale GRIDs layouts covering at once CPU needs, transmission needs as well as storage needs. They may be called **WAG** (for Wide Area Grids)

Key questions to day to be addressed to day are two-folds:

- How practically ensure a smooth and easy running of GRID-based system
- How to practically take advantage of GRID properties to implement a new vision for computerized system development and implementation of storage needs?

4) AN EXAMPLE OF VIRTUAL DATA STORAGE

In a conventional data base concept, the activities are logically organized as follow:

- Locate the information (catalogues etc..)
- Get hold of the data (media, access costs, transmission & delays)
- Process the information
- Transfer the result to suitable place




Fig.4 : A Wide Area Grid concept

In a WAG data based concept it should lead to :

- Create a Grid connection with the potential databases
- Make their archive accessible from this GRID (without being obliged to know in advanced what will be needed)
- Enable the creation of jobs submitted to the GRID

In fact there is no need anymore to exactly locate where the data stands or whether they are or not dispersed. Once a database is “seen” from the surrounding GRID it is “seen” from any user application jobs and the handling of communications and data repositioning should be handled by the “Upperware” itself.

=>>Then could be demonstrated the feasibility of a **Virtual Factory** able to dematerialize the handling and the storage of information

5) PRELIMINARY CONCLUSIONS

Virtual Storage is a logical consequence of GRID concept evolution. Of course large-scale database will still be needed and present archiving system will stay. But their insertion (or immersion) in the frame of GRID large scale topology (WAG) with the help of suitable virtualization as allowed by advanced operating systems for GRIDS (Upperwares) will induce a drastic change in the way to “think” our archive. No surprise Europe including ESA and EU will become key players in association with other major European space systems suppliers like CNES. As for GRID and Internet , the coming technological evolution will be based on a natural selection process where success will be mainly driven by users addiction. The advent of “fully virtual” processing and storage systems in which design may become as straightforward as writing a diagram on a piece of paper is to be encouraged since one of the key to lower cost and expand drastically the number of services offered to users. This can be considered as the coming Internet-like evolution for process management (as was Internet for information management). Needless to say most of this evolution will be “power-hungry “in term of needed resources for storage and communications ,inducing for sure, an other level of sizing for most of our archives. All these parameters are inter-dependants since no use to get hold of big archives if one cannot exchange or process them, (which sounds good for satellite telecomm prospect as well !). The advent of GRID and its upcoming evolution toward WAG and supporting improved Upperwares will create this synthesis. Obviously Europe by investing on such advanced topics may continue being at the forefront of computer system advanced evolution with obvious consequence in term of industrial prospect including “Space data information systems”. CNES is presently exploring many of these new options in the frame of the MERCATOR project.



OPERATIONAL SOLUTIONS FOR LONG-TERM PRESERVATION AND ADVANCED ACCESS SERVICES

Gilles PETITJEAN

*EADS Defence and Communications Systems
Intelligence Surveillance & Reconnaissance
Multi Mission Ground Segments
6, rue Dewoitine – BP14
78142 Vélizy-Villacoublay Cedex – France
gilles.petitjean@eads.com*

ABSTRACT

EADS Defence and Communications Systems (EADS DCS) has developed an expertise as integrator of archive management systems for both commercial and defence Customers (ESA, CNES, European Commission, EUMETSAT, European MOD, etc.).

The concern of valuable data owners is both their long-term preservation but also the integration of the archive in their information system with in particular an efficient access to the archived data for their User Community. EADS DCS, as system integrator, answers to this requirement by a methodology combining understanding of user needs, exhaustive knowledge of existing solutions, development and integration ability. EADS DCS has developed a set of off-the-shelf solutions, relying on CCSDS OAIS model and allowing an easy integration of powerful and reliable data servers.

This paper intends to analyse how the data owner can provide advanced access services to the archived data with a focus on the concern of long-term preservation requiring archived data migration to new storage hardware and software solutions.

The services offered to the User Community interested in having access to archived data have to take into account both the specificity of these data but also the needs of the Users. EADS DCS experience in the implementation of such services demonstrates the need for specific solutions tailored to each field of application. The search and selection of products archived by the Users require a catalogue entity presenting a user-oriented view to the archive content. The consistency between the archive and this catalogue is a key issue for the implementation of such archive facility.

The long-term preservation leads to data retrieved from the archive in a format and on support, which are not convenient to end-users. The data delivery to the users requires an additional component complementary to the archive in charge of reformatting, possibly compressing and disseminating the products either electronically or on media. The long-term preservation of archived data involves an organisation of the storage independent of physical supports, guarantying a stable access to the valuable data, even after migration of the archived data onto new storage media.

The paper presents EADS DCS experiences and generic off-the-shelf solutions for the implementation of a high-performance archive facility tailored to the specific needs of a User Community.

INTRODUCTION

EADS Defence and Communications Systems (EADS DCS) is an information system integrator for both commercial and defence Customers. The data managed by EADS DCS information systems are mainly images acquired by earth observation and meteorological satellites and by airborne systems (surveillance aircraft or unmanned aerial vehicles). Because these data have to be maintained available for long-term, EADS DCS has developed a comprehensive expertise as system integrator of archive management facilities (AMS/ADAR archive facility for ESA/ESRIN, U-MARF archive facility for EUMETSAT, HELIOS 1&2 archive for French MOD, SPOT 1-5 archives for CNES & Spot Image and for direct receiving stations, Eagle Vision Transportable Station for US DOD, etc.)

For EADS DCS Customers, the long-term preservation of their data sets is a key issue and some of them have already been confronted with the necessary migration of their valuable data to new storage technology. But it is not sufficient to design an archive management facility with the only concern to achieve long-term preservation. Why preserving for long-term the valuable data if their access by a designated User Community is impossible or very painful. The design of an archive management facility shall take into account both the long-term preservation and a complete set of access services tailored to these data.

CCSDS OAIS MODEL AND THE ACCESS TO ARCHIVED DATA

The Consultative Committee for Space Data Systems (CCSDS) has defined the Open Archival Information System (OAIS) reference model [1]. OAIS is a common framework of terms and concepts for the definition and the design of future archive systems facing with permanent or indefinite long-term preservation of digital information.




Fig. 1: Functional entities of the OAIS Reference Model

The access to archived data is supported by the Access functional entity. This entity can be broken down into two sub-entities:

- The “User Services” provide to the Consumer some information on the services offered for the access to archived data and supporting the search and order services to define the set of products to be prepared and delivered to the Consumer from the archived data.
- The “Formatting and Delivery” takes in charge the processing of the orders defined by the User Services. The selected data are retrieved from the archive and processed to the selected format, possibly compressed and delivered to the Consumer either by media or electronically.

The figure below presents the mapping of these two sub-entities with respect to the OAIS Reference Model.




Fig.2: Data Consumer Access of the OAIS Reference Model can be broken down into two entities.

The “Data Management”, also named “Catalogue” in various fields of application, shall present a consistent view of the “Archival Storage” content (also named here “Archive”), because the Consumer Access is based on selection of archived data from the Catalogue. The “Ingest” entity, in interface with the Data Producers for ingesting their data, is in charge of ensuring the coherency of both Archive and Catalogue contents.

USER ACCESS TO ARCHIVED DATA

The User Services are the interface between the owner of valuable data and his User Community. This interface is usually integrated in the on-line services offered by the data owner to his User Community. On the other side, the User Services are strongly depending on the nature of archived data and on the use of these data by the User Community. For these two reasons, it is difficult to envisage a generic solution for both User Services and Catalogue covering all types of valuable data and whatever the data owner organisation. The two following examples demonstrate this difficulty to envisage such a generic solution for various fields of application.




Fig.3: Example of User Services offered by U-MARF facility [6] to the EUMETSAT User Community for selecting meteorological data from geostationary or from polar satellites. The selection of archived data is done according to multi-criteria searches such as type of product, geographical area, range of time.



Fig.4: Another example of User Services offered by EUDOR facility [5] for OPOCE (EC) to the European Citizen for selecting and ordering legal documents issued by the European Commission. In this case, the selection of archived data is done according to key words present in the title or in the document itself.

Nevertheless, even if a common solution is not easy to define and design, EADS DCS experience has identified some common features between the various implementations.




Fig.5: Typical architecture for the implementation of the User Services and generic features

HOW OFFERING TO USER A RELIABLE VIEW TO ARCHIVED DATA

A usual problem for the owners of archived data is to maintain a consistent view of the “Archival Storage” content, considering that the Consumer Access is based on the selection of archived data from the Catalogue. In the OAIS Reference Model, the entity in charge of ensuring the coherency of both Archive and Catalogue contents is the “Ingest” entity.

EADS DCS has designed a generic solution, named “Generic Front End” or GFE [3]. Based on extensible architecture, GFE performs automatically all types of operations required during ingestion: quality control, reformatting or processing, metadata and browse generation, consistent cataloguing and archiving. This software handles heterogeneous data and performs parallel data ingestion. Originally developed for earth observation data, it is fully operational for other types of data.




Fig.6: The Generic Front-End (GFE) ingests concurrent data from data Producers and updates in a coherent manner both the Archive with the data and the Catalogue with the metadata describing the archived data.

GFE is designed as a plugin-based architecture. Plugins are used both for interfacing the data producers with tailored communication protocols (e.g. CORBA, SOAP) and for the processing of ingested data (quality control, metadata derivation, browse generation). The GFE customisation for each new data type is done by addition of new roadmaps (described as XML files) defining the sequence of plugins to be applied, including possible new plugins. The integration of existing processes is easy using standard Java™ mechanisms.




Fig.7: GFE is designed to manage heterogeneous data set in a configurable manner

The GFE component is already in use in UMARF facility [6] for ingesting all the meteorological data (METEOSAT) for EUMETSAT, in EUDOR facility [5] for ingesting all official documentation of EC for OPOCE (Office de Publication Officielle de la Commission Européenne). Moreover, GFE is now the reference ingestion solution for ESA in the Multi Mission Facility Infrastructure (MMFI) [8] in charge of handling all ESA Earth Observation data.

HOW DELIVERING ARCHIVED DATA TO USER

Once the Consumer has defined his order with a list of products to be delivered from selected archived data, the Access entity of the OAIS Reference Model has to process it.

As for the ingestion, EADS DCS has designed for the order processing and delivery a generic solution, named “Product Formatting & Delivery” or PFD [4]. Based on extensible architecture, PFD performs automatically all types of operations required for archived product delivery in co-operation with the Archive: retrieval from archive, possible reformatting, compression and dissemination, either physically on media (tapes, CDs, DVDs...) or electronically via network.




Fig. 7: PFD is a set of components managing heterogeneous data delivery from the archive to the end user in co-operation with the archive component based on the EADS DCS archive solution named SatSTORE

The main objectives for archived data delivery addressed by the combination of both PFD and SatSTORE solutions are the following ones:

- **Product formatting before delivery:** when defining his order, the User describes the sub-part of the data he is ordering and possibly a data format compatible with his working environment. Because the delivery formats can be multiple and can evolve during the lifetime of the archived data, this is a key issue to be able to propose various types of formats and level of processing independently of the archiving format. Moreover if the ordered product corresponds to a subset of the archived data, the ability to do this subset as soon as possible in the delivery chain can dramatically improve the performances especially for large sized data.
- **Delivery on heterogeneous supports:** multiple supports can be necessary for the delivery of archived data; whereas the electronic delivery meets a large success with the improvement of wide area networks, the media delivery remains useful especially for large sized data. Delivery on tapes (DAT, DLT, ...) are still in use for huge volume of data to be delivered, but most of the media deliveries are now done on CDs and DVDs.
- **Automatic delivery and operations:** because of the development of large archives using automated storage libraries, the access to the data is now mainly unmanned. This key evolution of the archive facilities allows data owner providing archived data delivery to user with reduced operational costs. In particular, the extensive use of electronic delivery authorizes fully automatic access services with low production costs. In case of media delivery, the use of robotics for media production (CD, DVD, tapes) optimises the resources of ordering production.
- **Security:** the electronic delivery requires secure mechanisms to allow giving access to the ordered products while the archive remains impregnable.

The PFD component is already in use in U-MARF facility [6] for delivering huge volume of meteorological data (METEOSAT) for EUMETSAT, in EUDOR facility [5] for delivering all official documentation of EC for OPOCE. Moreover, PFD is now the reference delivery solution for ESA in the Multi Mission Facility Infrastructure (MMFI) [8] in charge of delivering all ESA Earth Observation data and is already deployed in seven ESA sites.

ACCESS TO ARCHIVED DATA AND LONG TERM PRESERVATION

The pace of technology evolution is causing some hardware and software systems to become obsolete in a matter of few years. This is a key factor to consider when designing an archive management facility.

The architecture of the archive management facility shall allow separating the descriptive information used by the consumers for querying archived data and the archived data themselves as recommended in the OAIS reference model [1]. Moreover, the archived data must be consistent in term of application point of view and the storage software shall propose access services independently of the data storage location.

EADS DCS has developed the SatSTORE [2] solution to support an efficient integration of the archive management facility with the customer application by providing advanced services.

- Application oriented storage organization: the interest of the concept of archived products instead of a simple file handling is that all the components of a product (image, metadata, browses, other) are archived together under an application-oriented archive product identification.




Fig. 5: SatSTORE Product concept separates the client view on archived data and their internal storage organisation

- Data migration: data have to be migrated when a storage technology (hardware and software) becomes obsolete or when its cost of ownership (in particular, maintenance costs and storage volume in term of square metres) justifies the migration to a new one. SatSTORE proposes several mechanisms to manage this migration in a transparent manner for the other components of the OAIS model (Ingest and Access entities). In particular, the separation of the internal storage organisation and the access services to archived data guaranties the access to data in case of long-term preservation. Moreover, the ability to reformat the archived data before delivery allows envisaging an evolution of the archive format of the data during their lifetime independently to the delivery formats.

SatSTORE solution is the reference solution for major European organisations such as ESA, EUMETSAT, EC, French MOD. SatSTORE is currently in operation in about 100 worldwide configurations with heterogeneous capacities and performances (from few to several hundreds of terabytes).

REFERENCES

- [1] *Reference Model for an Open Archival Information System (OAIS)*
CCSDS 650.0-B-1 Blue Book January 2002
- [2] *SatSTORE - High Capacity Archiving System*
EADS Defence and Communication Systems documentation
- [3] *GFE – Generic Front-End Overview*
EADS Defence and Communication Systems documentation
- [4] *PFD – Product Formatting & Delivery Overview*
EADS Defence and Communication Systems documentation
- [5] *EUDOR web site – The portal to European Law*
<http://europa.eu.int/eur-lex>
- [6] *Unified Meteorological Archive & Retrieval Facility (U-MARF) Requirements Specifications*
EUMETSAT document – EUM/UMA/SPE/001 Issue 2 Rev.1
- [7] *Archive Management System (AMS) for ESRIN EO Data – Software Requirement Document*
ESA document – AMS/LR/SRD/001 Issue 2.0
- [8] *ADAR, a Study for an Advanced Data ARchiving System for Earth Observation*
White paper presented to SpaceOps 2004 in Montréal -By G.M. Pinna and P. Mougnaud from ESA, G. Petitjean from EADS, C. Demange from GAEL Consultant

EUMETSAT Unified Archive and Retrieval Facility: Status and Plans

F. Cadé
EUMETSAT
Am Kavallerisesand 31
D-64295 Darmsadt, Germany

ABSTRACT: The EUMETSAT Unified Meteorological Archiving and Retrieval Facility (U-MARF) will be the central repository for all EUMETSAT data and centrally derived products. It will also provide users with an off-line retrieval capability using web-based interfaces. The first version of the Facility was developed in the frame of the Meteosat Second Generation Programme (MSG) in order to archive all EUMETSAT geostationary products, up to the latest satellite, MSG-1 now operationally known as METEOSAT-8. A second version of the Facility is currently under development in the frame of the EUMETSAT Polar System (EPS), in order to extend the services to the low-earth orbit products and data flows, and deal with the significant planned growth of the user community. This paper presents the plans for the development of the U-MARF. It describes the evolutions currently undertaken, as well as all the capabilities and services that are ultimately expected from this Facility.

1. INTRODUCTION

The radiometer from the Metosat Transition Programme (MTP) satellites, currently under operation in EUMETSAT, provides 1.8 Gbytes of Level 1.5 data every day in three spectral bands, culminating to more than 650 Gbytes per year. Due to the enhanced spectral and temporal capabilities of the SEVIRI instrument flying on-board the Meteosat Second Generation (MSG) satellites, about 60 Gbytes of data and Meteorological products are generated daily since the launch of MSG-1 in 2002, now operationally known as METEOSAT-8. This represents more than 21 Tbytes of data per year and per satellite. With the development of the new EUMETSAT Polar System (EPS) Programme this figure will rise even more drastically to 88 Tbytes per year and per satellite, whilst the EUMETSAT user community is also expected to widen significantly. Moreover, in the context of the scientific exploitation of these data and products the needs for off-line retrieval services are planned to grow exponentially, and issues related to delivery media, storage cost or batch processing should not be underestimated.

As part of the evolution from a single mission operator to a multi-mission operator, and in order to face the challenge of growth in archiving and off-line retrieval needs, EUMETSAT has analysed the commonality among its Programmes. One outcome was the decision taken by EUMETSAT Council in 1997 to develop a multi-mission, incremental archiving infrastructure, that includes a web-based user access capability: the Unified Meteorological Archiving and Retrieval Facility (U-MARF) that will support all EUMETSAT Programmes, i.e.:

- The Meteosat Transition Programme (MTP), including historical archive from METEOSAT 1 to 7,
- The Meteosat Second Generation (MSG),
- The EUMETSAT Polar System (EPS).

This Facility should also be able to extend to future EUMETSAT Programmes, as well as Third-Party missions data and products.

2. DEVELOPMENT APPROACH AND STATUS

2.1. Project Drivers and objectives

The decision to build a common archiving infrastructure for all existing and on-going Programmes was also driven by the following considerations:

- The development of the user community across the various Programmes and the increased awareness of the potential value of historic meteorological data, especially in light of the new EUMETSAT mission to support meteorological and climate requirements from its Member States;
- The need to offer the user an integrated access to all EUMETSAT products and off-line services, i.e. non-real-time services, across all Programmes;
- The rationalisation of development and operational costs by promoting the use of Commercial-Off-The-Shelf (COTS) items and industrial projects re-use wherever possible;
- The opportunity to take advantage of new technologies (including archiving, on-line web technologies) and emerging international standards (such as the CEOS concepts);
- The need to define an infrastructure that builds in flexibility and scalability to accommodate future growth in data, archive and retrieval requirements.

2.2 Development approach

The facility is being developed following an incremental life cycle around two major versions.

The development of the first version, the U-MARF V1, was phased with the MSG Programme development. It aimed at providing common services such as the user administration and specific services related to the Geostationary (GEO) data and products from MTP and MSG. Due to the late start of the U-MARF project with respect to the MSG Programme, the V1 development consisted in its turn of two major sub-versions:

- V1-A version allowing to ingest, catalogue and archive data, and a limited monitoring and control, and therefore capable to secure the MSG programme by supporting the most critical pre-integration activities with the MSG Ground Segment,
- V1-B containing all hardware and functions, and especially the complete retrieval functions and user services.

The second version, the U-MARF V2, is being developed in parallel with the EPS Programme. It will extend the existing services to the EUMETSAT Polar System (EPS) data and products. It also provides some inter-operability with the Archiving and Retrieval services offered by the Satellite Application Facilities (SAF) installed across EUMETSAT Member States.

Similarly the development of the V2 is being performed with two major subversions V2-A and V2-B according to the same breakdown of functions as for the V1: the V2-A aiming at supporting pre-integration activities with the EPS Ground Segment, and the V2-B containing all hardware and functions.

Both versions of the U-MARF Facility are procured from industry based on a contract awarded in November 1998 to a Consortium lead by EADS / Defence and Security Systems, and involving Advanced Computer Systems (ACS), as a major sub-contractor for the development of the Catalogue and User Services.

2.3 Development status and Plans

- **V1 status**

The industrial development of the V1 version was kicked off in December 1998. The Provisional Acceptance was completed in March 2002. EUMETSAT proceeded then with the integration of the Facility with the MSG Ground Segment. The Facility started ingesting, cataloguing and archiving MSG-1 data and products in November 2002 in the frame of the MSG-1 commissioning. The hand-over to the Operations was formally announced in January 2004 as the general outcome of the MSG-1 Readiness Operation Review.

In terms of availability of services, the whole ingestion chain (i.e. ingestion of data until archiving and cataloguing) was fully operational for the MSG-1 launch. The standard web-based user services will be fully operational, after connection to the secure EUMETSAT infrastructure.

- **V2 evolutions**

The development of the V2 version was kicked off in April 2002. The detailed design phase was completed in June 2003, with some delay in order to rectify design issues and accommodate evolutions of the EPS technical baseline. The V2 development is now running according to plan. The Provisional Acceptance of the V2-A version was achieved in March 2004, and this version is being used in order to support investigations and pre-integration activities with the first version of the EPS Core Ground Segment.

The development of the full V2-B version is also on schedule. The V2-B production phase was completed in July 2004, whereas the authorisation to start the Integration and System Verification Phase already given May 2004. The Provisional Acceptance of the V2-B version is currently planned at the end of January 2005. Beyond this point, additional versions are to be developed in order to accommodate evolutions required by the EPS Programme and the SAFs, while performing the integration and system verification with the EPS System.

3. U-MARF TECHNICAL OVERVIEW

3.1. U-MARF Functions

- **General Presentation**

The U-MARF consists of 3 main functional chains:

- the ingestion, archiving and cataloguing chain,
- the retrieval chain,
- the monitoring and control functions.

The ingestion archiving and cataloguing chain consists of specific front-end subsystems in charge of managing the interface with the individual ground segments, a standard front-end in charge of common functions enabling the ingestion of granules into the U-MARF, the archive subsystem which manages the archive itself, and finally the catalogue subsystem which manages the storage of the metadata of all the archived granules. The synchronisation and the consistency between the operations of archiving and cataloguing, is also ensured by the standard front-end.

The retrieval chain consists of the User Services, the Catalogue and the archive subsystems in the retrieval mode, and finally the Product Formatting and delivery in charge of delivering the products to the end-users.

The monitoring and control functions are ensured by a specific subsystem. However, the scheduling of the retrieval requests, and the administration of the users is part of the User Services subsystem.

Such an architecture was defined in order to ensure a maximum flexibility with respect to future evolutions of the Facility, and in particular in order to accommodate the ingestion and management of new sets of data and products from additional ground segments from future Programmes.

The figure 1 below presents a technical overview of the U-MARF with its decomposition into major subsystems. Also shown on the figure, but outside the U-MARF domain, are the current ground segments for MTP, MSG and EPS, from which the U-MARF Specific Front Ends receive data once the satellites are operational.




Figure 1 – U-MARF System Overview

The function and purpose of these subsystems are described briefly hereafter:

- **MSG Front End (MFE)** : This subsystem receives telemetry, images and meteorological products data from the MSG ground segment (namely the IMPF and MPEF systems respectively). It checks these products for consistency and completeness, extracts metadata and in some cases creates quick-look images. It then submits the metadata and the U-MARF granules containing the data and quicklook images to the GFE.
- **MTP Front End (TFE)**: This subsystem receives image and meteorological product data, and in some cases browse images, to be migrated from the MTP MARF. It extracts the metadata and then submits the metadata and the U-MARF archive granules, containing the MTP data and browse images to the GFE.
- **EPS Front End (EFE)**: This subsystem receives multi-sensor data from the EPS ground segment. It checks the data for consistency and completeness, extracts metadata and in some cases creates a browse image and the related geolocation data. It then submits the metadata and U-MARF archive granules, containing the data and browse images to the GFE subsystem.

- **SAF Front End (RFE):** This subsystem receives SAF metadata and browse images from the SAFs. The validity of the metadata is checked and along with the browse images, is submitted to the GFE. This subsystem has been designed in order to be able to ingest also products from the SAFs.
- **Standard Front End (GFE):** This subsystem collects products in U-MARF standardised U-MARF archive granules from the various front ends (MFE, TFE, EFE, RFE), performs further granule and metadata consistency checks, generates browses (from the MSG quicklooks) and thumbnails. Here the operator can intervene if necessary. It then supervises the correct cataloguing and archiving of the products within the catalogue (CAT) and archive subsystem (STO).
- **Archive subsystem (STO):** The task of this subsystem is to manage the U-MARF archive, ensuring a secure long-term archiving of all U-MARF granules after ingestion, and supporting all data retrieval requests in response to user orders. It also performs the administration of the archive, including management of the organisation of the archive, and administration of the archive media. The subsystem shall support the capability to migrate the archive media to a future new technology.
- **Catalogue Subsystem (CAT):** This main function of this subsystem is the management of product metadata. It supports the ingestion process by storing the product attributes of each new archived product, supports retrieval requests by passing all product catalogue/attribute information to the User Services to allow it to serve user search and order requests, and performs the management of the metadata data base management system.
- **Product Formatting and Delivery subsystem (PFD):** This subsystem processes data retrieved from STO (following user orders) into the format requested by the user, after applying possible spatial and/or band subsetting options. Processed products ready for delivery are then either written to media or placed in the FTP Delivery Space according to the user-specified delivery method.

The subsystem consists of an entire infrastructure of peripherals, and controls all additional operations required for the preparation of the shipment of the media to the end-users (including the labelling of the media, the shipping notes...).

In order to provide a maximum flexibility, the re-formatting software are defined as set of plug-ins modules that comply to an internal interface.

- **Monitoring and Control (UCM):** The UCM subsystem provides the tools to monitor the “health” of the U-MARF by monitoring various hardware and software parameters, and if necessary generate reports, and to control the U-MARF by the ability to start and stop subsystems as necessary, and change operational modes.
- **User Services subsystem (US):** This subsystem aims at providing a complete set of web based services in order to access to U-MARF data and products. These services are basically organised into:
 - core U-MARF services, including information and guide, product search (inventory and browse functions), ordering and follow-up,
 - supporting auxiliary services including users registration and administration, help desk.

The subsystem can be decomposed into three components running on different platforms. The User Services Administration (USA) runs on the catalogue server interacting with the catalogue data layer and is in charge of processing the orders and the management of the overall subsystem. The User Services Interface (USI) runs on the HTTP server and provides access to the Users for searching, ordering and retrieving products. Thirdly, the User Services Client (USC) is downloaded directly to the Users' computer to facilitate the interaction of the User with the U-MARF by means of a Java Applet providing an intuitive tool for the specification of queries and visualisation of results.

Beyond the implementation of the services themselves, the subsystem is in charge of the complete user administration (USA). This includes implementation of the EUM Data Policy, and of the management and follow-up of the whole ordering process, from the receipt of the user's electronic order request to the point at which products are ready for delivery and billing information sent to EUM Finance Division for invoicing.

- **U-MARF V2 evolutions**

The MTP Front-End, MSG Front-End, Archive Management, Catalogue Management, User Services, Product Formatting & Delivery and Monitoring & Control subsystems have been developed for the V1 version.

In the frame of the V2 development, two new specific subsystems are being further developed: the EPS Front-End, and the SAF Remote Front-End, as described above.

In addition, a major redevelopment of the Catalogue and User Services subsystems is being undertaken in order to introduce the EPS and SAF specific services, and to provide some inter-operability with the SAF Catalogue and User Services. The User Services Client layer was also completely revisited with the introduction of Java Applet technology.

The design of the Product and Formatting subsystem was slightly modified with additional peripherals, and new reformatting software being developed as plug-ins for the additional EPS products and external distribution formats required for the V2 version, in particular with the introduction of HDF5 format.

Last but not least the physical architecture has been upgraded in order to migrate from a DEC Alpha architecture under OSF, to an architecture based on SUN Solaris, and the initial archive library changed in order to move to a final StorageTek robot that will provide more expansion capability. This migration was undertaken on both V1 and V2 versions, respectively by EUM Operations Department and by the U-MARF Contractor.

3.2. U-MARF Services

- **General Presentation**

The U-MARF has been designed to provide users with automated delivery services and to offer the maximum flexibility in terms of selection of delivery media, spectral bands, geographical areas, and selection of repeat cycles within time periods. The U-MARF standard services are described in the U-MARF User Service Handbook. These services will be accessible from the EUMETSAT web portal (c.f. EUM home page www.eumetsat.de).

The U-MARF Services can be basically classified into U-MARF core services and auxiliary services.

The core services include:

- Information and Guide services aiming at providing users with general information and documents (about the U-MARF services, the products characteristics, the file format, etc.);
- Product search services allowing users to search for archived products and to obtain information on the selected products (these services correspond to the Inventory and Browse services from the CEOS Catalogue definition);
- Ordering and Follow-up services allowing users to order the previously selected products.

The auxiliary services support the aforementioned core services and include User Registration and User Administration, as well as a Help Desk service to complement all the automated services listed above.

The U-MARF retrieval function will be fulfilled by the Product Search and Ordering services. The search services aims at providing user-friendly access to the data catalogue and the archived products. Several types of searching mechanisms are proposed:

- Simple product search limited to time range, satellite identification, product type and overall product quality,
- Advanced product search adding more specific criteria such as ingestion time, missing line count, quality flag...
- Standing Order, a more sophisticated search definition with a repeat time period.

In response to a query, a user will obtain the list of requested products, their main characteristics (metadata) as well as thumbnails of the related products. Users can also have access to browse image data prior to proceeding with ordering.

The ordering services will allow users to order the selected products. Several ordering options will be proposed in order to allow a refinement of the user order:

- Subsetting: selection of specific spectral bands, definition of a geographical region of interest;
- Format definition: native format (band interleaved format for MSG Level 1.5 data, or WMO format for MSG Level 2), band sequential (for MSG Level 1.5 data), JPEG;
- Compression method;
- Delivery definition: on-line or physical media (paper, CD-ROM, DVD, DAT, Exabyte, DLT).

• U-MARF V2 evolutions

The overall set of the U-MARF services as described above remains the same, except for the introduction of new simple and advanced search for the LEO products services, the SAF product services and the multi-mission services.

In order to manage the search for the orbit-based LEO products and some SAF products, the concept of a region of interest search is introduced that allows the user to request only the products which intersect a chosen region of interest. This region of interest is defined by its numeric co-ordinates values, or by a graphical tool directly on a map. In the answer to these user requests, the swath of the orbits matching the request is also displayed in addition to the list of the products characteristics and their thumbnails.

A multimission search is also added that allows the user to search for data, satisfying a particular search criteria, across more than one mission (GEO, LEO, SAF).

Finally, beyond these different search services directly based on satellites, instruments and products, a more conceptual thematic search will be put in place that will allow users to search across a collection of products.

The User Services Client layer was also completely revisited with the introduction of Java Applet technology necessary for the implementation of the new orbit based search, and in order to improve the overall ergonomics and look and feel, with similar principles as the one developed by ESA.

3.3. Performances issues

Meeting performance requirements have always been the major challenge of the Facility.

The archiving capacity results from the average ingestion throughputs from the various ground segments and the operational usage of the Facility. The U-MARF is expected to be operational until 2020 with an overall storage capacity up to 1700 Tbytes. The archiving profile currently planned is depicted in the Figure 2 after.




Figure 2 – U-MARF archive profile

With respects to throughput requirements, the U-MARF V1 was designed to ingest data from the MTP archive (28 Gbytes /day) and from 3 MSG satellites (180 Gbytes /day), and to provide a nominal daily retrieval capability of 100 Gbytes per day.

For the V2 with the introduction of LEO satellite data and SAF metadata, the requirements are for the ingestion on the V1 load (MTP and MSG), plus data from 2 Metop and 2 NOAA satellites (480 Gbytes /day), SAF metadata and browse images (100 Mbytes /day), together with a retrieval capacity of 600 Gbytes per day.

In order to accommodate the potential increase in the number of users over the operational lifetime of the U-MARF the system is expected to have a growth potential to handle a daily retrieval capability up to 1.1 Tbytes /day.

Finally, in view of the operational usage of the Facility, stringent availability requirements apply. In particular, a requirement for a Mean Time To repair of 10 minutes, lead to the concept of definition of a nominal and back-up ingestion chain, allowing to switch rapidly in case of any blocking problem, or planned unavailability of the archive or of the catalogue.

4. CONCLUSIONS

The development of the U-MARF facility is part of EUMETSAT’s effort to build an incremental multi-mission infrastructure, exploiting commonality between the different EUMETSAT programmes.

The general objective is to provide a secure long-term archiving, answering the growing demand of EUM user community, and to enhance/foster the user access to EUMETSAT products and services as they become relevant to a wider range of user communities in meteorological, climate and scientific fields.

Understanding Indexing-Contextures with a View to Add Value on Models and Datasets in e-Science

Ying Liu⁽¹⁾, Fen Yang⁽¹⁾, Anita Coleman⁽²⁾

⁽¹⁾*Computing Science Department, University of Aberdeen
Kings College, Scotland, UK
{yliu,fyang}@csd.abdn.ac.uk*

⁽²⁾*School of Information Resources & Library Science
University of Arizona
Tucson, AZ 85719 USA
asc@u.arizona.edu*

INTRODUCTION

One of the main objectives in e-science is to establish a problem solving environment (PSE) that provides the computational facilities needed to solve a target class of problems [3]. These features include advanced solution methods, automatic and semiautomatic selection of datasets, solution methods, and ways to easily incorporate novel solution methods. Preserving datasets to provide an added value product and service in PSE is desirable to scientific communities. We argue that we not only need an adequate indexing approach to the preservation, but also must incorporate views of the datasets' usage and subsequent transformations.

The concept of indexing is not new [38]. Almost 50 years ago, Fairthorne claimed: “indexing is the basic problem as well as the costliest bottleneck of information retrieval” [30]. In his mind, the fundamental indexing-idea was that one may use computers working with (or on) a “text” and producing from that work an “index entry” representing the text. To date, the text and the meaning are still two fundamental abstractions. The former is more concrete, but still varies a lot. The latter is more abstract, and is often hard to “pin down”.




Fig.1 Contextures of Scientific, syntactic, computational and semantic communications in problem solving environments

What challenges us today is a set of dynamic indexing contexts in PSE. Advances in technologies enable us to digitise, archive, visualise and publish datasets on the internet. This has created many new opportunities in organising datasets through, e.g., gateways [28] and repositories [7]; however, scientific datasets often require a retrieval that is coherent with a model for solving problems.

For example, in land-use management, land-cover images are fundamental and valuable data sources. The segments on such an image, e.g., an area of forest or grassland, are identified and classified by scientists who undertake a series of mathematical analysis and processes supported by the chosen algorithms and software tools. In order to understand or express these scientific association clearly we describe the relationships between the notion of “indexing”, “models” and “datasets”. We find that

- a) the indexing research in data storage structures such as H-tree, R-tree, or n-Dimensional space has common assumptions made for organizing arbitrary files [35], not for identification and classification of scientific models.
- b) currently the web based infrastructures, such as the metadata [26][27][11], the semantic web, ontologies, the web services, XML repositories [13], digital libraries [2][4][5][10], middleware, the semantic Grid, have enabled PSEs (see Fig. 1). These infrastructural technologies have enabled IR research for information discovery rather than collocation problems, e.g., distributed retrieval, or contextual, knowledge based, user-driven or hybrid retrieval [36].

In the rest of this paper we outline the essential indexing problems. We describe “bibliographic entities” (bib-entities) as a way to facilitate the indexing processes. We then address the potential applications of this approach by a “real world” case in the management of land-uses using artifacts from scientific communication and research such as models and datasets.

THE INDEX PROBLEMS

Distinguishing between “text” and “meaning” is precluded by the inflexibility and inadequacy of current approaches to organizing “datasets” and “models”. For true scientific re-use of the information retrieved, we need to know how the datasets are generated, why the datasets are collected, and how the datasets are used. IR (Information Retrieval) studies focus on known-item identification based on disciplinary subjects [12][34][36]. Datasets (some of which may be collected in real-time while others may be generated through modeling), detected segments in images, and models, can all be called *works*. Recently information scientists are exploring the cataloguing and indexing of scientific items by treating scientific models and datasets as *works*, and in turn, classifying and indexing the *works* as entities [24]. “A work is an abstract entity; there is no single material object one can point to as the work. We recognize the work through individual realization or expressions of the work, but the work itself exists only in the commonality of content between and among various expressions of the work” [6, p. 34]. A focus on *works* will help to advance understanding of the role *works* play in facilitating knowledge construction, and in the importance of the work entity as key to the construction of bibliographic and specific databases, or internet search engines.

Currently two major approaches to providing intellectual access to scholarly web resources are in place. A major initiative for applying metadata to the web resources was established as CORC (Cooperative Online Resource Catalogue), a joint program of OCLC and a cadre of participating libraries. Participating libraries select the web resources they believe to be most valuable and then catalogue them, including Dewey Decimal Classification numbers to provide subject access. The other approach is the creation of “portals” or gateways. Special libraries and information centres can provide an important service by identifying those web resources of greatest relevance and value to their users, indexing those resources, and developing a gateway that provides access to these resources through metadata elements. Many efforts are also a hybrid of both approaches to knowledge organization for information retrieval. New prototypes of domain-specific knowledge services are under research and development in line with the portal approach, e.g., Edinburgh GIS Portal, the Alexandria Digital Library, the THREDDs (the THematic Real-time Environment-

al Distributed Data Services), or DLSI (Digital Library Service Integration) projects [7]. However, it is imperative to systematically tackling the problems on the following four dimensions.

On the Bibliographical Dimension

The indexing databases focus on the subjects of disciplines (what topics and concepts are there within a particular subject or discipline) that result in products that can be indexed for IR. Subject analysis including classification is concerned much less with disciplinary, experimental and intellectual *works*. For example, bibliographic records of datasets on water runoff reveal little about objects such as land-area, water, hydrologic bodies, or processes, such as saturation and infiltration.

On the User-community Dimension

Our studies of scientific communities show that:

- a) online archiving scientific *works* is highly desirable;
- b) the indexing must cope with online-raw data which are semi-structured and file based;
- c) components of scientific *works* often are numerical models, formatted datasets, file-based, or operated instrumentally;
- d) components of scientific *works* are usually acquired or systematically derived from a scientific simulation, modeling environment, or data instruments;

On the Web Dimension

For knowledge services, segregating tools such as indexes or catalogues is insufficient. Scientific *works* (e.g., datasets) must be further represented as bib-entities, so as the properties of bib-entities would ground the adequate datasets in the problem solving processes. Furthermore, machine-enforced descriptive and subject catalogue-rules can then be formulated.

On Metadata Modelling

Metadata standards such as Dublin Core [26] are used as the content standards. But many corresponding encoding schemes are missing. Thus, the main contribution from a metadata model is limited to enable the indexing practice adding values of the objects, phenomenon, or processes as parts of different types of subjects. Thereby scientific datasets as *works* should be reinforced to hold descriptions about some critical parameters of datasets' preservation in a PSE (see, again, Fig.1.)

GROUNDING BIB-ENTITIES IN THE INDEXING PROCESSES

For a given problem solving environment, a bib-entity serves to explicitly encapsulate the relevant datasets derived from users' workflows. Take users' models for example. As Abbott observed in the hydro-informatics community [31]: "the numerical model now becomes the domain knowledge encapsulator, for it is this which now encapsulate all that is known about the physical system that can be taken as given in any particular situation. This encapsulator must then itself be capable of using all the data that are available and of processing them in such a manner and with such rapidity that they can efficiently assimilated ...". In database communities for example, modelling the relationships between heterogeneous datasets are based on either Entity-Relationship or object-oriented modelling methodology supporting a query [37], such as "where is the nearest restaurant in town A?", not a query meant to reflect *works*, e.g., "what are the most accurate parameters calibrating this model?".

Bib-entities vs. Objects

Another fundamental problem of representing *works* as bib-entities is that, certain categories of things or other perceptual entities (based on the first-hand experience of knowledge) are widely used by people, because they are natural groupings [15][25]. The identities of these entities are often identified by their coherent clusters of features that appear in a variety of contexts [1][9][16][23][39]. For example, it may not be necessary to define what physical objects are and what their relationships are [40], but simply to accept groupings of “plant, water, land”. Further employment of the groupings is in preference to arbitrary or even bizarre contexts, such as, “Alfalfa, - a kind of crop – needs water there” may make perfect sense. This leads us to study how we represent the bib-entities encapsulating the indexing contexts. First we compare the difference between the concept of object originated in software engineering and bib-entity on the following four categories. (the object as a *technical* notion is extensively reviewed in [32]; also see [18][22] regarding OO software engineering methodology.)

Data abstraction. The behaviour of an object is encapsulated in its methods. Methods are mechanisms that have access to and can change the state of object. Thus, an object type is described in terms of the form of its instances and the operations (methods) applicable to its instance variables. The instance variables, together with the methods, are called the properties of the object.

In a bib-entity, the notion of methods for changing the state of an object does not exist. Instead, the concept of a law defines the lawful “allowed” states. Laws are viewed as properties. The existence is a basic principle, not an arbitrary decision.

Independence. Independence implies that the only way an object can change its state is through the actions of its own methods.

In a bib-entity, there is no explicit mechanism for the creation or destruction of things.

Message passing. Objects can communicate only through message passing. A message can cause an object to “behave” by invoking a method of an object. Thus, objects may affect each other through message passing and, since objects are independent, this is the only way they can do it.

In a bib-entity, anything lawful can act upon one other.

Inheritance. Objects can be grouped together into classes. A class is a definition of an object type. All objects in a class have the same instance variables and methods and respond to the same messages. Objects can be categorised into subtypes or subclasses through specification. Specification forms an “IS-A” or “PART-OF” hierarchy. Objects in a subclass inherit all the instance variables and the methods. Multiple-inheritance can exist; therefore, class hierarchy can be viewed as a directed acyclic graph.

In a bib-entity, classes are not at all obvious. The concept of a class is allocated to a kind by considering the combined scope of a set of properties, and further to a natural kind of considering the scope of a set of rules. Thus, both “static” properties and “dynamic” are used in the definition of a kind of hierarchy. One concept or class may include another, e.g., “plant”, “crop”, and “Alfalfa”. A given entity may be assigned to two or more conceptual classes which are not hierarchical related, e.g., “river”, “land owner”, “pollutant”.

CASE STUDIES

In this section we illustrate how classified segments of land-cover images may better be organised as *works*.

The scientific background of land-use change is the proposed relationships between structured features of a landscape and its functions and processes. Lands have been converted to urban developments; forest,

woodland, grassland, and pasture have been converted to other uses ... These changes in land use have important implications for future changes in the Earth's climate and, in turn, great implications for the subsequent land-use change. Up-to-date knowledge of the state of the land and its dynamics are essential for making environmental strategic decisions.

Remotely sensed images are the main sources of raw data about the changing land- cover. The land-cover classes consist of vegetation types (such as forest, wetlands, and agricultural crops or pasture) and categories of non-vegetated surfaces (such as water bodies, bedrock outcrops, or settlements). Scientific in-depth analysis of the images for the segmentation and classification of earth observation not only results in huge amount of data to be handled, but also *works* (models and datasets) to be organised. Segmentation is the process of partitioning an image into segments, i.e., sets of adjacent pixels, having a meaning in the real world. For land-cover mapping, the segments should represent different land-cover categories. Classification denotes the process of assigning class labels to objects on the basis of a set of features. The objects may be individual pixels for segments. The process of segmentation may imply a change of representation from the digital image data structure (pixel raster) to a list and/or graph data structure. Classification (matching) is then performed on the objects represented in the list (graph). The process of covering this list of classified objects into a thematic map may again involve a change of representation from the list data structure back to an image data structure. Fig 2 shows the workflows.




Fig. 2. Computational segmentation and classification of land-cover images

A case is taken from [33]. The data in the experiment is from aerial images covering the Elgin area in north-east Scotland. The image has been manually segmented into different land cover types. The bibliographic entities together with the indexing codes representing the land cover type of the pixel are showing in Table 3. These six kinds of land cover types, i.e. arable land, good rough grassland, poor rough grassland, bracken, mixed woodland, and scrub are showing that land cover types are ecologically “similar”. Yet, classifying their differences is significant to land-use management, e.g., it can tell us that arable land likely transfers to good rough grassland and possibly transfers to poor rough grassland, bracken, mixed woodland, and scrub. The classified segments have to be organised as *works*.

Table 3. Land-cover types description

Land cover category	Land cover code	Main feature
Arable field	100	Arable field (no rock, no farms, no trees)
	150	Smooth grass/rushes(no rock, no trees)
	151	Smooth grass/rushes(no rock, trees)
Good rough grassland	155	Smooth grass/low scrub(no rock, no trees)
	156	Smooth grass/low scrub(no rock, trees)
	160	Undifferentiated smooth grass(no rock, no trees)
	161	Undifferentiated smooth grass(no rock, trees)
Poor rough grassland	140	Undifferentiated Nardus/Molinia(no rock, no trees)
Bracken	170	Undifferentiated bracken(no rock, no trees)
Mixed woodland	79	Undifferentiated mixed woodland(area)
Scrub	82	Undifferentiated low scrub

The used software tools and algorithmic methods are:

- a) Gabor filter as the discriminant classifier which makes use of the homogeneous texture descriptor TD derived from filtered images.
- b) MPEG-7 texture descriptors for the retrieval of images.
- c) Gabor-MPEG-7 integrator is a set of Gabor filters.

The classified segments are showing in Fig.4 & 5.




Fig. 4. Left polygon is labelled as arable field, and right polygon is labelled as smooth grass/rushes but no rock and no trees.




Fig. 5. An example of polygons in training data set. From left to right and up to down, the land cover code of these polygons are: 100, 150, 151, 155, 156, 160, 161, 140, 170, 79, and 82 as denoted in Table 4.

CONCLUSION

There is a growing need for collocation and retrieval of scientific datasets as *works* in a problem solving environment, not merely for the services on known-item identification based on subjects of disciplines. For scientific users, online archiving referenced scientific data is highly desirable. The indexing has multi-levels of semantic spaces and has to cope with online-raw data that are semi-structured or file based. We need bibliographic oriented entities that encapsulate datasets as *works*. The entities should also encapsulate observation or experiment abstractions and measurements within a given problem solving environment. Currently we are lacking of an adequate systematic approach to model such an entity.

REFERENCES

1. Albano, A., Bergamini, R., Ghelli, G., Orsini, R.: An Object Data Model with Roles. In: Agrawal, R., Baker, S., & Bell, D. (eds.): 19th International Conference on Very Large Data Bases. Dublin (1993) 39-51
2. Borgman, C. L.: What Are Digital Libraries? Competing Visions. *Information Processing and Management* 35 (3) January (1999) 227-243
3. Carpenter, B.: What Is Grid Computing? ISOC Member Briefings. (Grid Computing #11). 15 April (2003) <http://www.isoc.org/briefings/011/index.html>
4. Chen, S-S.: Digital Libraries - the Life Cycle of Information. BE (Better Earth) Publisher (1998)
5. Chen, X., Kim, D-h, Nnadi, N., Shah, H., Shrivastava, P., Bieber, M., Wu, Y-F.: Digital Library Service Integration. In: Proc. of the 2003 Joint Conference on Digital Libraries (2003)
<http://csdl.computer.org/comp/proceedings/jcdl/2003/1939/00/19390384.pdf>
6. Coleman, A. S.: "Scientific Models as Works." *Cataloging and Classification Quarterly* (2002) 33 (3/4): 129-159. Preprint available: <http://dlist.sir.arizona.edu/archive/00000060/>
7. [DLSI homepage]: Digital Library Service Integration. <http://www.is.njit.edu/dlsi/>
8. Dretske, F. I.: *Knowledge and the Flow of Information*, Basil Blackwell Publisher (1981)
9. Fayad, M. E., Johnson, R. E., *Domain-Specific Application Frameworks*, John Wiley & Sons (2000)
10. Greenstein, D., Thorin, S. E.: *The Digital Library: a Biography*, Washington: Digital Library Federation and Council on Library and Information Resources, 2nd Edition December (2002)
<http://www.clir.org/pubs/reports/pub109/pub109.pdf>
11. [Metadata homepage]: Information/Data/Metadata Management - General Resources.
http://www.cens.ucla.edu/Education/Documents/Information_Data_Metadata%20Mgmt%20Resources.htm
12. Morrissey, F.: Introduction to a Semiotic of Scientific Meaning, and its Implications for Access to Scientific Works on the Web. *Cataloging and Classification Quarterly* fall (2002)
<http://phoenix.liunet.edu/~smiragli/entities.html>
13. Nambiar, U., Bressan, S., Lee, M. L., Li, Y-G.: Current Approaches to XML Management. *IEEE Internet Computing*. July-August (2002)
14. [NSDL homepage]: Building a 'Memory' for the National Science Digital Library. Online (7) 2 January (2003) <http://www.npaci.edu/online/v7.2/index.html>
15. Ossher, H., Tarr, P.: Multi-dimensional Separation of Concerns and the Hyperspace Approach. In: Proc. Symp. Software Architectures and Component Technology: The State of the Art in Software Development. Kluwer (2001)
16. Papazoglou, M., Hoppenbrouwers, J.: Contextualizing the Information Space in Federated Digital Libraries. In: *SIGMOD Record* 28 (1) March (1999) 40-46 <http://infolab.uvt.nl/people/hoppie/papers/sigmod-record-1999.pdf>
17. Houstis, E., Gallopoulos, E., Bramley, R., Rice, J. (eds.): Problem-Solving Environments for Computational Science. Special Issue. *IEEE Computational Science & Engineering*. July/September (1997)
18. Schmidt, D., Stal, M., Rohnert, H., Buschmann, F.: *Pattern-Oriented Software Architecture. Volume 2, Patterns for Concurrent and Networked Objects*. John Wiley & Sons (2000)
19. Smith, T. R., Zeng, M. L., and ADEPT Knowledge Team, "Structured Models of Scientific Concepts for Organizing, Accessing, and Using Learning Materials,
http://esm.ucsb.edu/fac_staff/fac/frew/cv/pubs/2002_concepts.pdf
20. Star, S. L.: Got Infrastructure? How Standards, Categories and Other Aspects of Infrastructure Influence Communication. The 2nd Social Study of IT Workshop at the LSE ICT and Globalization, London, 22-23 April (2002)
21. Fonseca, F., Egenhofer, M., Agouris, P., Camara, C.: Using Ontologies for Integration

- Geographic Information Systems. Transaction in GIS 6(3) (2002) 1-33
22. Voelter, M., Schmid, A., Wolff, E.: Server Component Patterns. [Book homepage]
<http://www.servercomponentpatterns.org>
23. Wong, R., Chau, H., Lochovsky, F.: A Data Model and Semantics of Objects with Dynamic Roles. In: Jackson, M, Pu, C. (eds.): Proc. of 13th International Conference on Data Engineering. Birmingham, UK (1997) 402-411
24. Smiraglia, R. (editor): Works as Entities for Information Retrieval. Medford, Conn.: Scarecrow Press, 2002.
25. Zimmerman, D. W.: Could Extended Objects be Made Out of Simple Parts? An Argument for 'Atomless Gunk'. Philosophy and Phenomenological Research 51 (1996) 1-29
26. Dublin Core Metadata Element Set, <http://www.dublincore.org/documents/dces/>
27. Z39.50 Document, <http://lcweb.loc.gov/z3950/agency/document.html>
28. Koch, T.: Quality-controlled Subject Gateways: Definitions, Typologies, Empirical Overview. The International Journal of Digital Information Research and Use. 24 (1) 2000. 26-45 <http://www.mcb.co.uk/oir.htm>
29. Metadata for Hydrological Coverage, <http://edc.usgs.gov/sast/meta/fnl/hydro.html>
30. Fairthorne, R. A.: Automatic Retrieval of Recorded Information, Computer Journal, 1(1), 1958, 36-41
31. Abbott, M. B.: Hydroinformatics, Information Technology and the Aquatic Environments, Avebury Technical, 1991
32. Snyder, A.: The Essence of Objects – concepts and terms, IEEE Software 10(1), 1993, 31-42
33. Yang, F. & Lishman, R.: Land Cover Change Detection Using Gabor Filter Texture, in Proc. of Texture 2003, Nice, France. 2003, 113-118
34. Baeza-Yates, R & Ribeiro-Neto, B.: Modern Information Retrieval, Addison Wesley, 1999
35. Freeston, M., A General Solution of the n-Dimensional B-tree Problem, in Proc. Of SIGMOD Conf, San Jose, May 1995
36. Allan, J. (eds.): Challenges in Information Retrieval and Language Modeling, Report of a workshop held at the center for intelligent information retrieval, University of Massachusetts Amherst, Sep. 2002
37. - , Objects, XML and Databases, Information and Software Technology, 44 (2002), 195-
38. Lancaster, F. W.: Do Indexing and Abstracting Have a Future?, Annales de Documentation, 6 (2003) 137-144
39. Theodorakis, M., Anlyti, A.: Constantopoulos, P. & Spyrtos, N., A Theory of Contexts in Information Bases, Information Systems, 27 (2002) 151-191
40. Fabrikant, S. I. & Bittenfield, B. P.: Formalizing Semantic Spaces for Information Access, Annals of the American Geographies, 9 (2) (2001) 263-280

Session Added Value Services

“USE OF BROWSE OR PREVIEW IMAGES IN SUPPORT OF DATA ACCESS (PV-2004)”

5-7 October 2004, ESA/ESRIN, Frascati, Italy

**S. W. Doescher⁽¹⁾, U. S. Geological Survey
R. H. Sunne⁽²⁾, Science Applications International Corporation
M. Neiers⁽³⁾, Science Applications International Corporation**

**⁽¹⁾U.S. Geological Survey
EROS Data Center
Mundt Federal Building
Sioux Falls, SD 57198 USA
doescher@usgs.gov**

**⁽²⁾Science Applications International Corporation
EROS Data Center
Mundt Federal Building
Sioux Falls, SD 57198
sunne@usgs.gov**

**⁽³⁾ Science Applications International Corporation
EROS Data Center
Mundt Federal Building
Sioux Falls, SD 57198
neiers@usgs.gov**

ABSTRACT

For more than 30 years, the U. S. Geological Survey (USGS) Earth Resources Observation Systems (EROS) Data Center (EDC), Sioux Falls, SD, has provided access to multiple terabytes of remotely sensed imagery of the earth's surface. With the advent of the Internet nearly 15 years ago, preview browse imagery has been used in conjunction with other search techniques to support researchers determining the best available data for their purposes. This paper reviews the original purposes of browse and the methods used to determine the best choices for browse generation, representing a variety of the remotely sensed imagery. Historically, the Committee on Earth Observation Satellites (CEOS) community, made up of civil agencies heavily involved in earth observation activities, has reported on the exploitation and use of browse. The rapid advancement in networking, computation, and storage technology has not only prompted changes to the online methods to access and use the browse or preview imagery, but also has prompted an evolution of the browse concept. A survey of the CEOS members will explore their changes to the use of browse.

INTRODUCTION

For more than 30 years, the U. S. Geological Survey (USGS) Earth Resources Observation Systems (EROS) Data Center (EDC), Sioux Falls, SD, has provided access to multiple terabytes of remotely sensed imagery of the earth's surface. Early in the 1990s, the advent of the Internet and online systems created a need for browse (or preview imagery) to facilitate the quality and locational review of the archival holding of remotely sensed data. This paper provides a historical context and explores the current methods associated with browse or preview imagery. Through numerous technical discussions and brainstorming sessions, a general consensus and philosophy on the purpose of browse (or preview imagery) was formulated.

STRATEGY/PHILOSOPHY

The early purpose of browse was to provide a convenient method for online users to determine approximate feature coverage of selected products and to provide visual quality information as a means to determine usability of a specific product. The general technical strategy for browse was prepared independently of devices used to display and provided

browse in industry standard formats compatible with Commercial Off The Shelf (COTS) viewers. The targeted size of browse imagery was around 0.5 (MB), and was produced by reduction in spatial resolution and band representation (if relevant). Browse generation was planned to minimize data handling and use of resources, and reference or location related information was incorporated in browse to maximize stand-alone use with COTS viewing software. Reasons for browse from image sources included the following: determining extent and location of clouds; observing quality problems such as speckling, line drops, and sun glint; and confirming geographical location.

EARLY DEFINITIONS

The initial focus of browse at the USGS/EROS Data Center was directed on three data sets: Advanced Very High Resolution Radiometer (AVHRR), Landsat Multi-Spectral Scanner (MSS), and Landsat Thematic Mapper (TM). During early discussions on AVHRR browse, several options were considered. One browse option, consisting of a single band of the imagery, was created using every 4th line and every 5th sample, whereby the original 10-bit data were reduced to 8-bit. Band 2 was used for daytime data, and band 4 was used for nighttime data. For this option, an image of 5400 lines and 2048 samples was reduced to 1250 lines and 408 samples generated in a browse that was 0.51 MB in size. Another option considered was created using a multi-band browse. The data reduction considered was the same as the single band above, but in this case a browse was for all bands. This option of multiple bands provided a "custom" browse that was responsive to a user's request. Potential "custom" examples included the following: color browse (uses recipe RGB:2,1,1); normalized difference index or greenness image; Shark Classes (Sea, Sun Glint, Land, Cloud, Snow/Ice); individual bands; and potential other user defined models. This option generated a browse of 2.5 MB. This early selection of the AVHRR browse for a single band was driven primarily by technology considerations associated with the size of the browse image, which affected both the storage capabilities and network bandwidth speeds. The process of determining how to package the browse for delivery to the user followed the selection of a browse image.

RADIOMETRIC ADJUSTMENT

A radiometric adjustment typically is applied to remotely sensed imagery because the dynamic range of the image is narrow and stretching adjustments can improve the visibility for the user. Initially, the browse was stored without stretching, but stretch points that were previously determined and stored in the browse header information were used to image a fast radiometric correction at the time of delivery. The original intent was to enable the user to dynamically apply a radiometric adjustment to the browse. After several years, it was determined that users seldom applied the dynamic radiometric adjustment feature, and for simplicity the browse is now stored with an applied linear stretch. The browse stretch points are determined by using the histogram of the browse image. The lower stretch point ranged between the points of 2.5 percent and 97.25 percent of the accumulated histogram. Joint Photographic Experts Group (JPEG) compression of inherently low contrast data should not occur until a contrast stretch is applied. The JPEG compression algorithm tends to map large areas of low contrast data together, which results in a "blotchy" appearance (fig. 1) adding an additional requirement for the user to stretch the image upon receipt of the browse.




Fig. 1. Example of the effects of JPEG compression on low-contrast images.

RESOLUTION REDUCTION

In this early example, the spatial reduction of every 4th line and 5th sample was accomplished by a simple subsampling process. Early consideration compared various "lossy" methods of compression with some users participating in the determination of "acceptable." JPEG, with a quality factor of 75, was selected as the compression method and delivery format.

In the 1995, Jamie Mistein, working with the USGS and the National Aeronautics and Space Administration (NASA), refined a process for resolution reduction and improved the browse image characteristics [1]. The improvement of the wavelet process over the subsample method is documented in Tab. 1. The wavelet processing uses a convolution filter that preserves the high frequencies or edges as the filter is passed over the image. A separate pass of the wavelet process was required for a halving of the resolution. Because a resolution reduction of eight-fold was desired, three passes of the wavelet process were needed. In addition, each pass of the wavelet process requires additional processing time.

Recent experimentation with a "pnmscale" [2], publicly available programming code produced comparable results to the wavelet process, but requires less processing time than the three passes of the wavelet process. Tab. 1 characterizes the time to perform the resolution reduction on a Pentium III, with a 800 megahertz processor and 256 MB megabytes of memory. The original image was a digital orthophoto that was 6620 by 7688 pixels with a resolution reduction of every 8th line and every 8th sample with a resultant browse of 778 by 961 pixels. The "pnmscale" programming code uses a method where the resultant pixel is a weighted average of the covered pixels. This method tends to replicate the human eye's function, as it moves farther away from an image. A recently released version of pamscale incorporates the pixel mixing capabilities of the original pnmscale and also provides additional options.

Table 1. Example methods for resolution reduction

Resolution reduction subsection			
Method	Sub sample	3 passes of wavelet	Pnmscale
Processor time	14 sec	31 sec	24 sec

OPPORTUNITIES FOR BROWSE REFINEMENT

In 1992, the EROS Data Center began operations of TM and MSS Archive Conversion System (TMACS), which migrated Landsat data from high-density to DCRSi Cassette Tapes (DCT) output. During this process, the original set of Landsat browse was prepared by subsampling the MSS at every 6th line and every 6th sample and the TM at every 16th line and every 16th sample. Presently, 12 years since the first transcription, the EROS Data Center is about to embark on another migration activity to ensure long-term preservation of the data. The Landsat Archive Conversion System (LACS) will transcribe DCTs to 9940B tape, which has a 200 gigabyte capacity. The browse is generated during this transcription. In this case, the browse prepared with the "pnmscale" processing will reduce the resolution for the MSS to every 4th line and every 4th sample and the TM to every 8th line and every 8th sample. The result produces a browse with a size increased by a factor of four and with improved resolution.

In both operational modes, TMACS and LACS, the reduced resolution processing applied to three bands enable the creation of an RGB color composite. These reduced resolution images are saved in case a refinement to the final browse preparation is desired. The reduced resolution image bands are stretched, composited, and JPEG compressed with a

quality factor of 75 to a generation formatted browse. A JPEG comment field is added with the appropriate metadata information.

FILM BROWSE

Also, available from the USGS/EROS Data Center are film-based products from the film archive. The requests for the film products have declined and some of the film media have begun to degrade. The current strategy is to stop the delivery of the film-based products and to provide digital products. Currently, high resolution to digital scanning is cost prohibitive. An interim strategy, at this time, is to generate medium and low resolution browse to provide for high resolution scanned products on-demand. The medium and low resolution products are produced using Kodak DCS ProSLR/n (13.5 MP) digital cameras. For a 5-by-5 inch film source, the medium resolution at 600 dots per inch is 13 MB for black and white (B/W) and 38 MB for color and the output format TIFF. The low resolution (or browse image) is reduced to a 72 dpi product and is approximately 400 kilobytes and is stored in a JPEG format. The scanned high resolution image products produced with a Zeiss SCAI and Leica photogrammetric scanners will have a variable spot size where the standard product is 1200 dpi. These products are 120 MB for B/W and 360 MB for color in a TIFF format. Fig. 2 shows the automated roll film digitizing system. With this system, 5 years is needed to digitize the film archive of 8.6 million frames. The initial priority is to scan the film products that have the highest risk of degradation.




Fig. 2. Automated roll film digitization system

CURRENT USES OF BROWSE

Traditionally, browse images are used to aid the user's search and selection of Earth Observing Data. Systems such as Earth Explorer, (<http://earthexplorer.usgs.gov>), enable the user to specify their desired characteristics and to search the metadata. The results of the search request contain a pointer to the browse to view for acceptability. Recently, some systems have evolved that are characterized as "browse first." These systems provide users with an image-based mechanism to exploit the browse more directly as the primary mode of guidance. One example of such a system is Glovis (<http://glovis.usgs.gov>); an example screen from Glovis is shown in Fig. 3. Fig. 4 graphically depicts the increased usage of the Landsat browse over the last several years.




Fig. 3. Example screen from Glovis




Fig. 4. Landsat Browse Usage

CEOS COMMUNITY CHARACTERIZATION OF BROWSE USAGE

Historically, the Committee on Earth Observation Satellites (CEOS) community, made up of civil agencies heavily involved in earth observation activities, has reported on the exploitation and use of browse. In 1999, the CEOS, Working Group on Information Systems and Services (WGISS), Browse Task Team prepared a Browse Guidance Document [3]. This document provides general guidelines and best practices for preparation of browse imagery. The following updated Tab. 2 depicts various browse characteristics available for remotely sensed data.

Table 2. CEOS browse characteristics available for remotely sensed data				
Agency	Satellite/ platform	Sensor/ instrument	Browse characteristics	URL
CCRS	Landsat 4&5	TM	400 by 258 image, subsampled 1:16, ground sample spacing = 480 m, usual band inclusion = 2,3,4, JPEG compression	ceocat.ccrs.nrcan.gc.ca
CCRS	Radarsat1	SAR	256 by 256 image, black and white, JPEG compression	ceocat.ccrs.nrcan.gc.ca

CCRS	SPOT 1-3	Panchromatic Multispectral	250 by 250 image, subsampled ~1:12 or 1:24, ground sample spacing = 240 m, usual band inclusion = 1,2,3 or pan, JPEG compression	ceocat.ccrs.nrcan.gc.ca
CNES/ NASA	TOPEX/ POSEIDON, Jason-1	Altimeters Topex and Poseidon	Ocean Surface Heights at 6 km, 0.5 deg, and 1 deg.	podaac-esip.jpl.nasa.gov/poet/
ESA/ESRIN	ENVISAT	AATSR	○ The AATSR Browse is a 3-color image product derived from the Level 1B product. ○ Coverage: Product stripe up to 500 km in across-track direction ○ Radiometric resolution: 0.1 km ○ Pixel spacing: 4 km	earth.esa.int
ESA/ESRIN	ENVISAT	ASAR	ASAR Image Mode browse: ○ ASAR product generated when the instrument is in Image Mode. ○ Coverage: Product stripe up to 4000 km by 56 -100 km in across-track direction ○ Radiometric resolution: Product ENL ~ 80. ○ Pixel spacing: 225 by 225 m ASAR alternating polarisation browse: ○ ASAR product generated when the instrument is in Alternating Polarisation Mode. ○ Coverage: Product stripe up to 4000 km by 56 -100 km in across-track direction ○ Radiometric resolution: Product ENL ~ 75 ○ Pixel spacing: 225 m by 225 m ASAR wide swath browse: ○ ASAR product generated when the instrument is in Wide Swath Mode. ○ Coverage: Product stripe up to 4000 km by 405 km in across-track direction ○ Radiometric resolution: Product ENL ~ 30 to 48 ○ Pixel spacing: 900 m by 900 m for WS ASAR Global monitoring mode browse: ○ ASAR product generated when the instrument is in Global Monitoring Mode. ○ Coverage: 405 km in across-track direction. Product stripe up to 40000 km ○ Radiometric resolution: Product ENL ~ 11 -15 ○ Pixel spacing: 1000 m by 1000 m for GM	earth.esa.int
ESA/ESRIN	ENVISAT	MERIS	Composite top of atmosphere browse product. MERIS browse products contain a subsampled set of selected RR radiometrically calibrated Level 1b products (15 bands). The browse product contains 3 selected spectral bands (red, blue and green) of 0 to 255 intensity levels for each band. ○ Coverage: product stripe up to 1150 km in across-track direction ○ Radiometric resolution: 15 $\mu W / (m^2.sr.nm)$ at 865 nm ○ Pixel spacing: 1 km by 1km at nadir	earth.esa.int
ESA/ESRIN	ERS	ATSR	Quick Look size when framed according to ATSR product extension: 210 sample x 256 lines image. Descending frames (most daytime): color composite ch 11micron inverted BT and 1.6 Refl. Ascending frames (most nighttime) ch 11micron inverted BT. Lat/Long grids and coastlines. subsampled 1:2, JPEG compression	earth.esa.int earth.esa.int/esapub.esrin.esa.it/eq/eq52/buon52.htm

ESA/ESRIN	ERS	SAR	500 by variable image, 8-look, 200 pixel spacing, JPEG compression, ground range	earth.esa.int
ESA/ESRIN	IRS-P3	MOS	384x384x24bits image, full resolution, JPEG compression	earth.esa.int
ESA/ESRIN	JERS	VNIR	1024 by variable image, composite 3,2,1 RGB, Level-0 subsampled 1:4, 1024xvariable, JPEG compression	earth.esa.int
ESA/ESRIN	Landsat	MSS	1000 by variable image, composite 6,5,4 RGB, Level-0 subsampled 1:3, JPEG compression	earth.esa.int
ESA/ESRIN	Landsat	TM	960 by variable image, composite 7,5,2 RGB, Level-0, subsampled 1:6, JPEG compression	earth.esa.int
ESA/ESRIN	SeaStar	SeaWiFS	418 by variable by 24bits image, subsampled at 4 Kms, JPEG compression	earth.esa.int
JAXA	ADEOS	AVM	512 by 512 by 8bits image, composite 4,3,2 as RGB, subsampled 1:9.77, JPEG compression	eus.eoc.nasda.go.jp
JAXA	ADEOS	AVP	512 by 512 by 8bits image, subsampled 1:19.53, JPEG compression	eus.eoc.nasda.go.jp
JAXA	ADEOS	OCT	2048 by 1024 by 8bits image, subsampled 1:2, JPEG compression	eus.eoc.nasda.go.jp
JAXA	ERS	AIM	800 by 800 by 8bits image, subsampled 1:8 horizontal 1:8.5 vertical, JPEG compression	eus.eoc.nasda.go.jp
JAXA	JERS	OVN	512 by 512 by 8bits image, subsampled 1:8 horizontal 1:6.05 vertical, JPEG compression	eus.eoc.nasda.go.jp
JAXA	JERS	SAR	750 by 750 by 8bits image, subsampled 1:8 horizontal 1:8.53 vertical, JPEG compression	eus.eoc.nasda.go.jp
JAXA	Landsat	MSS	512 by 498 image, composite 7,5,4 as RGB, subsampled 1:6.33 horizontal, 1:4.37 vertical, JPEG compression	eus.eoc.nasda.go.jp
JAXA	Landsat	TM	640 by 498 by 8bits image; descending composite 4,3,2 as RGB and subsampled 1:12.04 horizontal, 1:12.05 vertical; ascending band 6 and subsampled 1:3.01; JPEG compression	eus.eoc.nasda.go.jp
JAXA	MOS	MES	512 by 450 by 8bits image, composite 4,2,1 as RGB, subsampled 1:4, JPEG compression	eus.eoc.nasda.go.jp
JAXA	MOS	VTI	512 by variable by 8bits image, subsampled 1:1.95 horizontal 1:1.79 vertical, JPEG compression	eus.eoc.nasda.go.jp
JAXA	SPOT	HP	512 by 512 by 8bits image, subsampled 1:11.72, JPEG compression	eus.eoc.nasda.go.jp
JAXA	SPOT	HX	512 by 512 by 8bits image, subsampled 1:5.86, JPEG compression	eus.eoc.nasda.go.jp
JAXA	TRMM	PR	3960 by 880/8bits image, subsampled 1:9, JPEG compression	eus.eoc.nasda.go.jp
NASA	Aqua	AIRS	False color visible, brightness temperatures at 11 m, at 39-km pixel resolution	daac.gsfc.nasa.gov/www/gallery/global_browse/
NASA	Meteor 3 Nimbus 7 etc.	TOMS	Level 3 browse is a 1 degree latitude by 1 ¼ degree longitude grid full resolution (not subsampled), GIF compression	toms.gsfc.nasa.gov
NASA	Quik SCAT	SeaWinds	Surface winds at 25-km resolution	podaac-esip.jpl.nasa.gov/poet/
NASA	SeaStar	SeaWiFS	Level-1 browse is a subsampled (every other pixel, every other line) version of the band-8 raw radiance counts image. Level-2 browse is a subsampled (every other pixel, every other line) version of the chlorophyll a image. Level-3 browse is a subsampled (every 8th pixel, every 8th line) version of the SMI image array.	seawifs.gsfc.nasa.gov/cgibrs/seawifs_browse.pl
NASA	Terra	ASTER/TIR	700 by 830 image, composite 13,12,10 as RGB, full resolution, 114 kB	glovis.usgs.gov/

NASA	Terra	ASTER/VNI R	700 by 830 image, composite 3N,2,1 as RGB, resampled to 1/6 original size, 114 kB	glovis.usgs.gov/
NASA	Terra	MISR	Color images from red, green and blue bands for each of nine cameras reduced to 2.2km resolution	eosweb.larc.nasa.gov/MISRBR /
NASA	Terra and Aqua	MODIS	True color visible, brightness temperatures at 11 m, global grid at 39-km pixel resolution	daac.gsfc.nasa.gov/www/gallery/global_browse/
NASA	Terra and Aqua	MODIS	True color visible, surface temperatures, Vegetation index, Sea Ice, Active fire detection, 1-deg Global, and 10 by 10 deg tiles at 5km pixels resolution	landweb.nascom.nasa.gov/cgi-bin/browse/browse.cgi
NASA	TRMM	PR	3960 by 880 image, HDF RLE compression	lake.nascom.nasa.gov/data
NOAA/ NASA	AVHRR	NOAA-7, -9, -11, -14, -16 and -17	Sea Surface Temperatures at 9-km, 18km, and 0.5 deg	podaac-esip.jpl.nasa.gov/poet/
USGS/EDC	AVHRR	LAC	408 by 1000 image, B&W subsampled 4 th line, 5 th sample, single band (band 2 for day, band 4 for night), JPEG compression (Q=75), size 75 KB.	earthexplorer.usgs.gov
USGS/EDC	Declassified Satellite Imagery	Film	1400 by 115 image B&W, scanned at 50 or 100 dpi, JPEG compression (Q=75), size 36 KB	earthexplorer.usgs.gov
USGS/EDC	EO-1	ALI	157 by 435 image, composite 4,3,1 as RGB, pnmscale 0.125, JPEG compression (Q=75), size 23 KB	earthexplorer.usgs.gov
USGS/EDC	EO-1	Hyperion	64 by 838 image, composite 40,31,13 as RGB, pnmscale 0.25, JPEG compression (Q=75), size 19 KB	earthexplorer.usgs.gov
USGS/EDC	Landsat 1-5	MSS	390 by 590 image, composite 7,5,4 as RGB, each band reduced every 6th line and every 6th sample, JPEG compression (Q=75), size 60kB.	earthexplorer.usgs.gov
USGS/EDC	Landsat 4-5	TM	350 by 350 image, composite 5,4,3 as RGB, each band reduced every 16 th line and every 16 th sample, JPEG compression (Q=75), size 75kB.	earthexplorer.usgs.gov
USGS/EDC	Landsat 7	ETM+	825 by 750 image, composite 5,4,3 as RGB, 3 wavelet passes, JPEG compression (Q=90), size 185 KB.	earthexplorer.usgs.gov
USGS/EDC	Scanned Aircraft	Digital Orthophoto Quads	950 by 750 image, B&W images subsampled by 2 followed by 2 wavelet passes, JPEG compression (Q=75), size 150kB.	earthexplorer.usgs.gov
USGS/EDC	Shuttle	SIR-C	256 by 200 image, B&W, proprietary software, GIF, size 45 KB	edcdaac.usgs.gov/sir-c/

OBSERVATIONS

In the early years of browse, the balance of technology combined with user constraints limited the ability to distribute high resolution browse imagery. As technology and bandwidth improve, the delivery of better browse (or actual) data can be achieved. The ability to put terabytes of data online for immediate viewing and download is feasible. Even though seamless delivery systems are available the need for browse functionality is still anticipated to support access to the data that will remain offline.

REFERENCES

- [1] J Ellison , and J. Milstein, "Improved Reduction-Resolution Satellite Imagery" The Aerospace Corporation. Available from <http://ftpwww.gsfc.nasa.gov/IAS/handbook/pdfs/irrsi.pdf>, 1995.
- [2] Pamscale is available SourceForge.net Project: Netpbm - <http://netpbm.sourceforge.net/doc/pamscale.html>.
- [3] O. Ochiai, B. Burford "Browse Guidelines Document (BGD), Version 1.00," CEOS Working Group on Information Systems and Services Browse Task Team (BTT), August 1999.

SITOOLS : Enhanced Use of Laboratory Services and Data

T. Levoir⁽¹⁾, L. Carbonnaux⁽²⁾, C. Gougeard⁽²⁾

⁽¹⁾CNES

18, Avenue E.Belin
31401 Toulouse Cedex
France

Email: Thierry.Levoir@cnes.fr

⁽²⁾CAPGEMINI Toulouse

8, rue de Mesplé - BP1155
31036 Toulouse Cedex
France

Email : Laurent.Carbonnaux@capgemini.com, Christophe.Gougeard@capgemini.com

INTRODUCTION

For several years CNES has been developing generic software systems in order to set up Data Centers, in particular SIPAD (*Système d'Information, de Préservation et d'Accès aux Données*, i.e. Information System for Preserving and Accessing Data) and its successor SIPAD-NG (New Generation SIPAD) [1].

The SITOOLS project enhances this offer by providing access to existing services and data which need to remain in the laboratories, especially due to high upgradability requirements implying the presence of the people managing them, and to simplify access to systems located in different laboratories.

SITOOLS therefore aims to demonstrate the feasibility of creating a set of software components :

- installed in the scientific laboratories and managed by these laboratories,
- providing access to services and data previously inaccessible,
- with the final objective of creating a distributed data center.

The 5 Target Concepts

Each SITOOLS component must be designed to be :

- Light : easy installation and maintenance to allow management by the laboratory engineers,
- Portable : straightforward installation on the laboratory information systems,
- Generic : in order to access data from different experiments and disciplines,
- Modular : easy modification or replacement of each feature, flexible use of the pre-existing laboratory services and resources,
- Interoperable : to interface with other distributed services or systems.

COEXISTENCE WITH SIPAD-NG

It is important to understand how the SITOOLS objectives fit into those of SIPAD-NG [1].

SIPAD-NG is a complete operational system for preserving and accessing data which implements proven off-the-shelf products and a sophisticated data model, in order to meet the requirements of a data center with a long-term mission or to provide a particular entity with access to data from several projects.

SITOOLS is the fruit of an R&T action aimed at providing a set of independent, portable tools, based on OpenSource products, which can be used to build a data access system via Web Services and one or more client applications. The system is not designed to operate independently. It does not provide the global administration and management functions necessary for a long-term system. Its purpose is to provide fast, easy access to data distributed throughout several laboratories.

USERS

SITOOLS is designed for 2 types of user :

- the end user, the scientist who consults the final system set up to look for and download useful data. He/she will be referred to as the "user" in the remainder of this document.
- the manager, the entity (scientific laboratory, organization, etc.) which will set up and run the system to provide access to its data.

GENERAL ARCHITECTURE

SITOOLS is based on a concept of services, interconnected by a virtual Web Services bus, offering a global service which includes basic services and added-value services (AVS), and accessed by one or more client applications. To set up and then run this system, the manager must perform several phases : installation, parameterization, configuration, administration. The system may run locally, on a single machine, or be distributed over several machines or even several laboratories.




Fig 1 : General architecture : Diagram of the services software bus

SITOOLS must not be dedicated to a particular project. It has to be flexible enough to handle different type of data, different type of tools, different project scale ...

SITOOLS must therefore be highly adaptable :

- The information model must be adapted to each user project :
 - o Information model refers mainly to the data types managed. For example, the access criteria for earth observation data and solar or planetary data are completely different. Similarly, on the same subject, data sets (homogeneous data generally obtained from a single experiment with the same processing level) may also have access criteria which are different or which do not follow the same syntax. SITOOLS must be able to integrate them in a single instance and connect them via synonym management.
- The global system can be set up on one or more machines, one or more networks or one or more sites, as required.
- The system must be able to use the tools currently available in the laboratory (database, processing tools, etc.).

SITOOLS COMPONENTS

SITOOLS is broken down into services and into client applications accessing these services. SITOOLS initially includes one client application, but any project using the tools can develop another client application using the SITOOLS Services.

This breakdown means that the SITOOLS user project can adapt to project requirements :

- The user Man Machine Interface (MMI) is separate from the system core (built as services),
- Each project uses all or part of the services available.

This breakdown also makes possible the upgrade of a service or a client application within the system with no modification in the other components (for instance, to take into account new requirements or technology changes).

The following diagram illustrates the software architecture of a SITOOLS instance implementing the various components proposed by SITOOLS. Each component is described in the remainder of this paragraph.




Fig 2 : Software architecture diagram

BASIC SERVICES

SITOOLS defines 5 types of basic service which can be "plugged in" to the SITOOLS virtual bus :

- On-line catalog : provides access to metadata and data accessible on-line :
 - o A catalog contains one or more data sets.
 - o A data set is characterized by a common structure (a single table), referenced in the catalog meta model.
 - o Each item of metadata is characterized by its type, its class as well as display and criteria-definition indicators.
 - o A specific class metadata item references the external data as an http url.
- This solution has been preferred to an "element/attribute" model, which is more dynamic but not as easy to maintain.

- Off-line catalog : identical to the On-Line Catalog Service with batch access to data.
 - o The Off-line catalog is implemented in the same way as the On-line catalog.
 - o Since the data is not on line, the catalog supplies the data items in batch mode when requested by the command processing service.
- Repository : links all accessible services to form a global system; it is the nerve center of SITOOLS, managing its global configuration :
 - o Catalog management and synonym association (providing the link between metadata with the same meaning, but not the same syntax for different data sets),
 - o Data set grouping management (for a given subject, for example).
 - o Management of added-value services, access, dynamic parameters.
 - o Management of user rights : rights to catalogs, services.
 - o Based on a configuration using XML files.
- Command Processing : end-to-end management of user commands,
 - o It handles the data command process.
 - o It is used to call batch type added-value services.
 - o Based on a state/transition controller set up at CNES : RIO[3]. It supports WebService called (SOAP with attachment) :
 - The command processing module is defined by a state diagram.
 - This diagram (see Fig. 3) is written in a specific language used to define UML 1.4 state diagrams, which then generate the Java code for RIO [3].
 - The Command Processing service proposes two services :
 - New command, with the following parameters :
 - o The list of urls of the data to be processed.
 - o The added-value service (AVS) used during processing, if any.
 - o The parameters for this service (delivery site, etc.).
 - Reception of data from the added-value service called, with the following parameters:
 - o The list of result files.




Fig 3 : Example of Command Processing state diagram

- User Space : manages the command and processing result space, where the user will find command results :
 - Use of the WebDAV (<http://www.webdav.org/>) standard for user space implementation.
(*WebDAV stands for "Web-based Distributed Authoring and Versioning". It is a set of extensions to the HTTP protocol which allows users to collaboratively edit and manage files on remote web servers.*)
 - The advantage is that metadata can be allocated to the deposited files.
 - Based on the Apache DAV module, which manages the rights.
 - Jakarta Slide is used as the Java client interface.
 - The user space is integrated in the SITOOLS client application using "taglibs" with the same interface as for the catalog (the user will not see any difference between the data previously requested and the data available on line via catalogs).
(*JSP taglibs define declarative, modular functionality that can be reused by any JSP page.*
<http://java.sun.com/products/jsp/taglibraries/>)
 - The Apache DAV module can be visible directly in the Windows Explorer; there are other DAV-specific explorers.

A service type is characterized only by the interface to this service. SITOOLS will supply some "instances", but the laboratories can create others using different technologies.

ADDED-VALUE SERVICES

The added-value services (AVSs) are independent software programs, interfacing with the system to provide new features (graphs, data-mining, 3D navigation, etc.). An AVS has no specific boundary. It may be an independent application, a SITOOLS services client application, a web or Windows/Linux application, etc.

Management of the AVSs in SITOOLS consists in

- defining the various types of operation of the AVSs and therefore knowing how to call them (call method, type, call parameters, etc.),
- connecting them with the information model elements (collections, data or browse sets), and therefore proposing them whenever useful to the user,
- managing access rights and user quotas,
- limiting the need of adding service types (a new service type involves significant modifications - repository, client applications, other used services, etc.).

There are 2 specific types of AVS, however, which have a clearly defined boundary :

- Batch Added-Value Service (B-AVS) : data processing AVSs during command execution :
 - Synchronous or asynchronous processing.
 - Receives the list of addresses of the files to be processed together with the useful parameters of the Command Processing service.
 - Retrieves the data on line or in the user space.
 - Carries out its processing, then calls the Command Processing service with the address(es) of the result file(s).
 - The Command Processing service retrieves the result and stores it, possibly zipped, in the user space.
 - It must implement a SOAP service interface which can be called by Command Processing. An http Form version is planned.
- Interactive Added-Value Service (I-AVS) : AVS used to work interactively on the user space :
 - Interactive processing with the user.
 - The I-AVSs can be internal or external :
 - Internal : standard function of the SITOOLS system, e.g. : selection aid AVS, thumbnail browser.
 - External : remote SOAP service or http Form called by the client application.

CLIENT APPLICATION

SITOOLS proposes a client application :

- Web application based on Struts for implementation of the MVC2 model. The MMI layer is in HTML 4 and CSS 2 generated in JSP.
- The super catalog (see Fig. 4) handles interaction with the catalog services authorized for a user.
- It dynamically prepares the display of the list of search criteria, integrating any synonym dictionaries.
- The client application supplies development APIs to access the various service types available.
- These APIs are Taglibs for JSP :
 - o Access to the super catalog service.
 - o Display of the metadata attributes according to display rank and class.
 - o Display during entry of search criteria according to display rank and class.




Fig 4 : Super catalog, dictionary of synonyms

USING THE TOOLS

The services connected to the SITOOLS bus together with the associated client applications form a global system providing access to data and services, i.e., a Data Center. This may be :

- a completely local system, on a single machine,
- a system distributed over several machines,
- a system distributed over several laboratories.

It may use one or more services proposed by SITOOLS and it may connect to existing services (database, AVS, etc.).

SITOOLS AND OPEN SOURCE

SITOOLS basic services are developed (or integrated) with Open Source tools exclusively.

Production tool :

- Apache Jakarta : for the http and DAV server.
- Jakarta Tomcat : for the application server part
- MySQL : for the DBMS system.
- Jakarta Struts : for the model-view-controller framework.
- JWSDP and Jakarta Axis : for the SOAP layer
- Jakarta Log4J : for logging management.

Development tools :

- Eclipse : Java IDE.
- JWSDP and Jakarta Axis : for the SOAP layer
- Jakarta Ant : for the deployment tools.

Since implementation complies with J2EE, JDBC and SOAP standards, it is quite possible to use other tools in production (Oracle DBMS, BEA WebLogic server, etc.)

CONCLUSION

SITOOLS is under development. The first version, V1, will be available in September 2004. It includes the main features required to validate the various architecture choices.

SITOOLS V1 must generate constructive feedback from users in order to improve the specifications of SITOOLS V2 which is to be developed for the summer of 2005.

To ensure they are representative, tests with SITOOLS V1 are being conducted by two scientific laboratories in France (IAS and CESR) as well as with MEDIAS, a company experienced in setting up data access systems.

REFERENCES

- [1] Heulet D. "SIPAD-NG : Multi-Field System for Accessing Scientific Data and Added-Value Services", PV 2004, ESA
- [2] Levoir T. (CNES Toulouse), Harvey C.C. (CNRS/CESR), Thieman J.R., Bell E.V. (NASA/GSFC), Huc C. (CNES Toulouse), "SPASE prototype : an example application for the interoperability in Space Physics", American Geophysical Union (AGU) Fall Meeting, San Francisco (USA), 8-12 December 2003
- [3] F. Jocteur-Monrozier, "'RIO Request Interface for Operations : A web service and workflow prototype", DASIA 2003, Prague

SIPAD-NG: a multi-field system for accessing scientific data and added-value services

D. Heulet⁽¹⁾, E. Houdet⁽²⁾

⁽¹⁾CNES
18, Avenue E.Belin
31401 Toulouse Cedex
France
e-mail: Dominique.Heulet@cnes.fr

⁽²⁾CS SI Toulouse
ZAC de la Grande Plaine
Rue Brindejonc des Moulinais
BP 5872
31506 Toulouse Cedex
France
e-mail: Emmanuel.Houdet@c-s.fr

INTRODUCTION

The SIPAD System

Towards the end of the 1990s, the CNES Plasma Physics Data Centre, known as the CDPP (Centre de Données de la Physique des Plasmas), developed a software package called **SIPAD** (Système d'Information, de Préservation et d'Accès aux Données), for enabling users to retrieve data and metadata [1].

The SIPAD system allows to:

- Catalogue data archived by the CDPP,
- Consult the catalogue via a Web server,
- Visualise browse products or documents related to the data through the Web server,
- Order data and receive them by different means (through networks or particular media),
- Implement added-value services (for instance: graphic displays of data ordered, extraction of fields from data files).

SIPAD was designed so that it could be used by other projects and for other fields than Plasma Physics.

Actually, in the past, CNES developed several systems to make scientific data from different projects available to users. These software packages were specially designed for specific scientific research projects, but nevertheless have many similar functions.

SIPAD has now demonstrated that it is possible to design software with all of the basic features found in any data access system, but which can be adapted to any scientific discipline.

SIPAD has thus been used to make systems for distributing data in the fields of **Plasma Physics** (for the CDPP, CLUSTER and CASSINI projects), for **Atmospheric Chemistry** (the ETHER project), for **Oceanography** (the MERCATOR project), for an overall cost lower than that which would have been necessary for separate specific development work for each of these projects.

The SIPAD-NG System

The **SIPAD-NG** (SIPAD – New Generation) system has benefited from five years of experimental results from the SIPAD system, the emergence of new needs among scientists and the potential offered by the most recent standards and technologies.

As for SIPAD, **the aim of SIPAD-NG is to avoid recurrent development** of data access systems in Mission Centres or Data Centres financed by CNES.

In particular, SIPAD-NG is intended to be used by thematic Data Centres whose long-term mission is to regularly enrich the list of data and added value services made available to a user community.

SIPAD-NG is thus an efficient, open-ended operational system with sophisticated administration functions.

FEEDBACK ON SIPAD

Over the five years in which the SIPAD system has been in operation we have drawn the following lessons:

- Generally speaking, a data access system is upgraded continually to meet new needs, to integrate new added-value services and to take into account changes in the computing environment,
- The Web User interface is, in particular, that part of the system which is improved through many releases in order to increase its user-friendliness and integrate added-value services,
- **In terms of advantages**, SIPAD offers an operational service; it has all of the basic features of a data access system; the information model defined for SIPAD has proved to be very robust and altogether suitable for the generic design principle; the data ingestion function (based on CCSDS recommendations for Data Dictionaries) is very open-ended and offers very stringent controls of the coherence and conformity of the information delivered,
- **Among its disadvantages**, the SIPAD system is not sufficiently modular (the METAPHASE product, which is used to manage the database, is omnipresent in the system); in spite of the improvements made, the response times for certain interactive functions are still too long (this is due to the internal workings of METAPHASE); moreover the Web User interface is not flexible enough to be adapted quickly.

Furthermore, new needs have emerged since the time when SIPAD's specifications were defined. These include:

- Needs for **added-value services**: even though SIPAD can to a certain extent integrate added-value services, it cannot meet the very demanding expectations of users in this respect. It is becoming more and more important to be able to provide data in a way that is suitable for users' needs, so that they can get the most out of the data. In particular, multi-sensor Data Centres now have to propose tools to compare data from different sources,
- Needs for **interoperability**: no Data Centre is expected to archive all data from all disciplines, or even data from only a single theme; interoperability enables a user connected to a Data Centre to obtain global information on the data which meet his needs, whether these data are available on the site being consulted or on another site. SIPAD does not have any interoperability function.

CHARACTERISTICS OF SIPAD-NG

SIPAD-NG's specifications were defined using feedback from SIPAD and from meetings between representatives of several projects from different scientific themes. A statement of work and technical specifications for SIPAD-NG were validated through a series of reviews involving the project teams.

Re-Using of SIPAD Functions

SIPAD-NG offers the basic SIPAD functions (for cataloguing data, for searching for and selecting data, for consulting information on data, for managing users and user profiles, for extracting archived data and for distributing data through networks or particular media).

SIPAD-NG uses same the Data Dictionary technology that allows to adapt the information model to match the characteristics of a particular scientific theme and to define and control the interface for supplying information to the system.

Integration of Added-Value Services

SIPAD-NG defines interfaces which can be used to integrate a wide range of added-value services. These include:

- Added-value services which can be used to process data before delivering them to a user,
- Added-value services which can be used to process data once they have been received in a user space.

An added-value service may have its own MMI, which can be used in particular for entering user parameters.

Interoperability

SIPAD-NG offers an interface to be consulted by a distant Data Centre. This interface is based on the 'Web Services' technology and on the Simple Object Access Protocol (SOAP).

Open-endedness of the User Web Interface

The SIPAD-NG has a Web interface for users to access data and added-value services.

This Web interface can be adapted in three ways to meet project needs:

- By changing its parameters: many parameters for appearance and composition of pages can be changed by simply initialising them in configuration files,
- Adaptation: some well-identified parts of the MMI JAVA code can easily be modified to adapt pages composition or linking,
- Specialisation: the Web user interface can be partially or completely rewritten without affecting the rest of the SIPAD-NG system. SIPAD-NG provides a **programming interface** for this purpose.

Modularity

SIPAD-NG has been designed with a highly modular architecture and, whenever the interfaces between components are not modified, it is possible to change a component without affecting the rest of the system.

GENERAL SIPAD-NG ARCHITECTURE




Fig. 1. General Architecture of SIPAD-NG

The SIPAD-NG kernel consists of:

- **Six basic services** which implement functions which can be found in any data access system:
 - The **'Catalogue Consultation' service** is used to search for data by navigating in data graphs and by applying selection criteria,
 - The **'User Management' service** is used to centralise all information about users profiles, access rights and quotas,
 - The **'Command Processing' service** is used to extract data from the archives and to possibly process them with added-value services before delivering them to users via network or media,
 - The **'User Space' service** offers users a workspace in which they can retrieve data and use added-value services,
 - The **'Ingestion' service** can be used to acquire update information on the SIPAD-NG catalogue, as well as information related to the data (browse products, documents),
 - The **'Administration and Supervision' service** is used to administer the system and to ensure that it is operating correctly.

- An **administration Web server** serves as the interface for the 'Administration and Supervision' service. This interface offers access to administration functions (add a user, monitor orders, etc ...),
- A **supervision synoptic** with an overall view of the way the system operates,
- A **database** which records global system information. This database is not the catalogue of data referenced by SIPAD-NG but only a base of parameters which are useful for the six basic services.

This kernel is an instance of SIPAD-NG. It is possible to install several independent SIPAD-NG **instances** on the same computer configuration.

A given SIPAD-NG instance can manage one or more projects.

As seen by SIPAD-NG, a **project** consists of:

- A **catalogue**: this is a database (used by the basic services) which references all of the information on:
 - Project data (the data organised into sets, which can be grouped into collections, the links between the data and the browse products or the documents). The data are not stored by SIPAD-NG; they can be accessed in an archive (which may be a near-line system or quite simply an on-line disk space) and the project catalogue memorises the location of the data in the archive,
 - Users (name, address, telephone number, preferences for delivery procedures, access rights to data, quotas),
 - Added-value services (links with data sets, quotas).
- **Disk spaces** for receiving:
 - Information for the ingestion function,
 - Browse products and documents corresponding to data,
 - User spaces,
 - Processing spaces for added-value services,
- One or more **client applications** which may be:
 - A **Web server** for the user interface: the project can use and adapt the client 'User Web MMI' application provided by SIPAD-NG; but he can also develop his own Web server using the programming interface to communicate with the basic services,
 - **Processing chains** or **added-value services**: the programming interface provided by SIPAD-NG can be used by any kind of application for ordering data. An application can issue a data order and then wait for these data to be supplied by SIPAD-NG,
 - **Internet applications**: an application located at a remote site (for instance another Data Centre) may consult the SIPAD-NG catalogue via Internet using the 'Web Services' technology (interface provided by SIPAD-NG).

In a given instance of SIPAD-NG which manages several projects, each project is completely independent of the others. The SIPAD-NG can be used to define project administrators: a project administrator has access to all of the administration functions for the project environment for which he is responsible (catalogue, disk spaces, client applications) but has no access rights to other projects.

For a SIPAD-NG installation at the CNES Computer Centre, a given instance will manage several projects. On the other hand, for a SIPAD-NG installation in a Data Centre outside of CNES, a SIPAD-NG instance will probably manage a single project.

With this architecture, it is possible to **set up several Web servers which are completely different in appearance**, thus enabling separate user communities to look for data according to different criteria and providing access to different added-value services, **while using the same SIPAD-NG services kernel**. The development effort then involves making the Web servers. It is not necessary to redevelop all of the mechanisms for managing the catalogue, users, orders, etc....

TECHNICAL CHOICES

The design of the basic services and MMI (user and administrator) for SIPAD-NG has a classic 5-layer architecture shown in Fig. 2.




Fig. 2. Logical Architecture for Basic Services and MMIs

Database

SIPAD-NG uses the **ORACLE 9i** Relational Database Management System for the instance database and project catalogues. ORACLE is efficient and widely-used in systems developed at CNES. The SIPAD-NG is intended to be used by major Data Centres which are able to purchase ORACLE licences, or have already done so.

Basic Services

The six basic services are developed with **JAVA**. They can be interfaced by a client application in two ways:

- By means of a **JAVA API** (used by an application installed on the same machine),
- By using the **RMI protocol** (used by an application installed on a different machine).

This means SIPAD-NG can be deployed on hardware with different architectures:

- **Single-machine installation:** the whole of SIPAD-NG and its client applications are installed on a single machine,
- **Multi-machine installation:** for instance, the Web server is installed on one machine and communicates via the RMI protocol with the basic services installed on another machine.

The call parameters for basic services and results sent back to client applications are **character strings in XML format**.

This technology means that it is no longer necessary to change the basic services methods signatures used by client applications. Any changes are introduced by adding elements in XML character strings transmitted between the basic services and the client applications, and it is no longer necessary to upgrade the existing client applications.

Client applications

SIPAD-NG has two Web servers which are specific client applications:

- **User Web MMI:** using the APACHE TOMCAT server and STRUTS framework,
- **Web Services interface:** using the APACHE TOMCAT server and AXIS for implementing the SOAP protocol.

Operating systems

SIPAD-NG has been developed for **SOLARIS** and **LINUX** operating systems.

AN EXAMPLE OF DEPLOYMENT

In the example shown in Fig. 3, the CDPP (Plasma Physics) and MERCATOR (Oceanography) projects use a single SIPAD-NG instance.

Each project has its own resources: data archive, catalogue, user Web server, administration Web server.

The Web servers have specific features to match their particular theme. For example, the MERCATOR server can be used to select the data according to geographic criteria, whereas the CDPP server cannot.

Moreover, the CDPP can inter-operate with other Data Centres which are part of the Space Physics Archive Search and Extract consortium (SPASE) [2].




Fig. 3. Two Projects sharing a SIPAD-NG Instance

CONCLUSION

2003 was devoted to specification studies for SIPAD-NG, which were validated by reviews involving the client project teams.

The **CS SI company** was entrusted with development which began in February 2004. A preliminary design review was held in June 2004 to validate the technical choices.

The SOLARIS version of SIPAD-NG has been scheduled for release in May 2005 and the LINUX version for release in July 2005.

The first SIPAD-NG client projects will be CDPP, MERCATOR and CASSINI, for which it will be necessary to migrate data from SIPAD in order to ensure the continuity of services offered by SIPAD.

The ICARE (clouds and aerosols) and AVISO (altimetry) Data Centres are potential clients for SIPAD-NG (contacts have been made and studies are being conducted).

REFERENCES

- [1] Nonon-Latapie M. (CNES Toulouse), Besson I. (CS SI Toulouse), Gangloff M. (CNRS/CESR), "The preservation and valorisation of scientific data by the 'Centre de Données de la Physique des Plasmas' (CDPP, Plasma Physics Data Centre)", PV 2002, CNES
- [2] Levoir T. (CNES Toulouse), Harvey C.C. (CNRS/CESR), Thieman J.R., Bell E.V. (NASA/GSFC), Huc C. (CNES Toulouse), "SPASE prototype: an example application for the interoperability in Space Physics", American Geophysical Union (AGU) Fall Meeting, San Francisco (USA), 8-12 Decembre 2003
- [3] Levoir T. (CNES Toulouse), Carbonnaux L., Gougéard C. (CAPGEMINI Toulouse), "SITOOLS: Enhanced Use of Laboratory Services and Data", PV 2004, ESA

Data description based on XML dictionaries

Arnaud LUCAS and Denis MINGUILLON
CNES data experts

1. INTRODUCTION

XML technology is more and more widely used.

XML is a very powerful syntax. But it is not enough to master a syntax to write meaningful things. Those who master French are not necessarily able to write like Victor Hugo.

Whatever the application domain is, XML technology has to be combined with a domain know how.

That is particularly true for the data description domain.

2. THE PROBLEM

Data had to be described before the arrival of XML. Many of the characteristics of data (ownership, volume, frequency, etc...) are intrinsic and independent from the underlying technology used to build the description. Several standardization efforts have led to the definition of normalized ways to describe data.

A very minimal core of attributes (common to almost any data) such as name, definition, comment, version etc... has been defined as the "Dublin Core" [ref 1] and is widely used specially in the library management domain. Another list of standard attributes has been defined by a CCSDS recommendation (CCSDS stands for Consultative Committee for Space Data Systems) named DEDSL standing for Data Entity Dictionary Specification Language (ref 2 & 3). This recommendation offers a standard way of extending the list of predefined attributes by creating user attributes to describe features not covered by those predefined attributes. (to name, definition, comment etc... one can so add other characteristics as volume, frequency, validity_range etc...).

So, the problem was to take advantage of XML technology without throwing out all this standardization work.

3. THE APPROACH

From the various standards tested, we learned that prior to jumping on XML it was important to list the characteristics that will be used in the construction of data descriptions. The same approach is common to any dictionary. For example in a French/French dictionary, for each item, there exist its name to designate the item, its nature (verb, adjective, name...), its definition and if necessary an example of its use.

Once this list is agreed, the dictionary can be built with respect to this list.

This approach has been formalized in a methodology procedure manual [ref 4] with no assumption about any standard.

This manual only contains a few main principles.

- to agree on the list of the characteristics
- then it encourages the separation between the content of a data description (the values of the description attributes belonging to the agreed list) and the presentation aspect.

XML is then the ideal tool to manage separately the content and the presentation.

4. A CONCRETE CASE IN A PROJECT

The approach had to be validated on a real project. The PLEIADES project was an opportunity.

The choice of XML for most of the interfaces between the subsystems had been made.


We had to convince the project team to instantiate the methodology manual for their practical case.

So, the first step was to define the list of attributes they wanted to see in their documentation.


Once this done, a tricky point was to map this list on an XML schema to be applied to all the XML files descriptions.

Some of the attributes matched naturally XML schema concepts. For some other ones an XML schema implementation had to be defined.

This step issued a project methodology document [ref 5] and a corresponding generic XML schema (see annex). A commercial XML tool (XML Spy) was then customized to help users to apply the schema to their XML data. The trick was to modify the internal Schema representation in the tool. We have used the fact that any kind of data could be inserted into the `xs:annotation` element, so we inserted our specific elements under this one.




Generated with XMLSpy Schema Editor www.xmlspy.com
Original annotation element [fig 1]



Generated with XMLSpy Schema Editor www.xmlspy.com
Modified annotation element [fig 2]

Then, we were able to use our XML Editor to support the edition of the interfaces (see fig 3), which were remaining compliant with the W3C standard.



XML Spy editor help with specific schema [fig 3]

The design of the various interfaces could then start.


5. PROJECT ORGANIZATION

In parallel of a methodology definition, this document defined a project organization (in the data description domain). This project manual promotes the usage of a centralized dictionary allowing sharing the data definitions needed by several sub-systems in their interface definition. This dictionary is made possible by the use of the formalized common schema. Someone responsible for a reference dictionary was designated. He is in charge of building the shared dictionary and promoting its use.

Each responsible (in the various sub-systems) in charge of some interface definition has access to the reference dictionary to build its definition in its local repository.

A way to propose local definitions to become part of the reference (to be shared with other sub-systems) has been established.

A survey across the various domains has been decided to avoid multiple definitions of the same data.



Project organization [fig 4]

The next step was to ingest existing dictionaries into the reference, not to reinvent the wheel. But each subsystem has its own standards which are in most of the case not compatibles.(ISO 19115/1930, ESA EOLI (earth observation data definitions), CNES internal dictionary...).The problem is still under discussion.

6. PROS & CONS

The main difficulty consists in the change in the work habits.

People are used to define data in a very informal way, using a non specialized tool as Word.

So they get the feeling that specialized tools are not convenient for the preliminary phases where data are only roughly known. So they are tempted to carry on with an informal definition they intend to translate later in a formal way when everything is precisely defined.

This is a big issue because the translation appears afterwards as a supplementary work. So the translation could be entrusted to someone else who starts from informal information with all the uncertainty it implies.

The acceptance of the recommended methodology needs a lot of explanation about the fact that formal and rigorous could fit in with preliminary and incomplete. For instance, it is easy to build a type named `TO_BE_DEFINED` and use it for all the concepts no to be further described at present.

Another issue is the fact that data description is always done in a hurry for the next project review that is coming soon. So the learning of new skills is always seen as a waste of time.

XML is maybe not convenient for all data. For data that must remain binary (images for in-

stance) there is a need to rely on domain standards if they exist. The link with related XML metadata has to be designed cautiously.

Finally, it is a fact that XML tools are not yet domain specialized. They do not yet capitalize a know how. They present XML rather as a goal than as a practical hidden support technique. This is changing and soon specialized tools will propose XML based output.

7. GAINS IN TERMS OF PRESERVATION

The first one is obviously that rigorous description will remain understandable. Informal description offers room to omissions and inaccuracy.

A second one is that encouraging the sharing of description causes a de facto standardization of the description of a domain. Users become familiar with the way of describing data. To build the PLEIADES reference dictionary a recommendation was to try to reuse as much as possible existing dictionaries.

Another gain comes from the choice of XML itself as the data syntax. XML is self describing by nature. It prevents from having data of which the separate description has been lost.

This is true if some rules are applied. Let's take the case of the unit associated to a measurement.

An XML schema can impose that the unit is mentioned as an attribute of a tag:

`<ANGLE unit="degree">`.

In that case data is self describing.

The unit can also be given in the XML Schema annotation only. In that case the danger of loss of the description remains.

One could say that all the description must be part of data but, then, there is a potential problem of volume to store.

8. CONCLUSION

The experience related by the current article is still relatively new at the time of writing: the PLEIADES project is still in its preliminary phase and data descriptions are not finalized.

The expected gains are more hope than a real fact.

The Earth Observation is a continuous activity that could take advantage from a stable description of the handled data. XML dictionaries could contribute to share a domain repository.

Reference documents :

[ref 1] - <http://dublincore.org/documents>

[ref 2] - DATA ENTITY DICTIONARY SPECIFICATION LANGUAGE (DEDSL)
ABSTRACT SYNTAX (CCSD0011)
CCSDS 647.1-B-1 BLUE BOOK
<http://www.ccsds.org/CCSDS/documents/647x1b1.pdf>


[ref 3] - DATA ENTITY DICTIONARY SPECIFICATION LANGUAGE (DEDSL)
XML/DTD SYNTAX (CCSD0013)
CCSDS 647.3-B-1 BLUE BOOK
<http://www.ccsds.org/CCSDS/documents/647x3b1.pdf>

[ref 4] - REGLES POUR L'ELABORATION DE DOCUMENTS
DE DESCRIPTION D' INTERFACES INFORMATIQUES
RNC-CNES-E-40-507

[ref 5] – Gestion des Interfaces Système Pleiades en XML
PHR-NT-0-1458-CNES

9. ANNEX

The schema of the metadata inserted into the xs:annotation element.



Infrastructures to improve user access to EO data

P. Seifert⁽¹⁾, S. Kiem⁽¹⁾, E. Mikusch⁽¹⁾, W. Willdegger⁽¹⁾, M. Böttcher⁽²⁾
Hans-Juergen Wolf⁽¹⁾, Holger Maass⁽¹⁾, Klaus-Dieter Missling⁽¹⁾

(1) German Aerospace Center (DLR), German Remote Sensing Data Center (DFD)
Oberpfaffenhofen, D-2234 Wessling, Germany
e-mail: Peter.Seifert@dlr.de, Klaus-Dieter.Missling@dlr.de
www: <http://www.caf.dlr.de>

(2) Werum Software & Systems AG, D-21337 Lüneburg, Germany,
www: <http://www.werum.de>

INTRODUCTION

A huge and permanently increasing amount of remote sensing data has to be archived and is supposed to be utilized. Therefore the classical issues of long term preservation of products of various missions and sensors have to be joined with easy and open data access. To cope with requirements of both aspects a separation of core systems and user information services (UIS) is indicated. This separation of duties generates flexibility in connecting multiple data sources with several user service systems.

The German Remote Sensing Data Center (DLR) has developed and is operating the Data Information and Management System (DIMS). The main component of the core part of DIMS is a digital library, the so-called Product Library. It is composed of an inventory part hosting the metadata of all remote sensing products and an archive part with a hierarchical storage management system for product files. The Product Library provides the product reference by a consistent metadata and data management as well as data security. Further core components of DIMS manage user orders, customer information, product order options and production workflows. Processing systems are connected to process and format products or to generate delivery products for online or offline distribution.

In comparison to the functionalities of the core components, the user information services provide the public access to the valuable products. In a user information service the emphasis lies on the user point of view to the data. Usually, higher level products are offered which can be used immediately without further processing steps. The products offered may exist as such in the core system, or they are producible on demand. This way additional products are offered which will be generated by a post-processing step performing tailoring of products originating from the core system. Therefore, the product collections in a user information service will be quite different to the ones in a core system. For external users the data source is of marginal interest, so the collections are sorted thematically. Quick search and user friendly ordering features are important issues in this domain.

The first part of this paper presents the bridge between the DIMS core systems and the user information services. Data from one Product Library has to be loaded to different portals for data access and also data from different Product Libraries can be provided in one user information service. With the UIS-to-core interface component of DIMS it is possible to exchange products and order/user information by appropriate mechanisms including trigger, timer and query. Upload of new products or product updates in one-by-one or bulk mode can be configured. User orders can be mapped to and from user information services of different organizations, e.g. ESA or EUMETSAT.

In the second part common, underlying hardware aspects are discussed to improve access to Earth Observation data. Future Earth Observation (EO) missions have also new challenges to the EO data management. Raw data and products are larger and new applications (real time, near real time) need the products in shorter time. But experiences of the last years show that a detached reflection on reduced access time to the archive will not lead automatically to shorten delay for the provision of user products. An integral approach to the whole production chain of user products enables the application of new storage architectures in the remote sensing settings which is necessary for future missions. After a short description and analysis of

the existing infrastructure the second part of this paper provides proposals of a set of guidelines and suggestions for the application of new storage architectures. These proposals will be subjected to some tests whose first results round off this lecture.

EARTH OBSERVATION PRODUCT MANAGEMENT: BRIDGING CORE SYSTEMS AND USER INFORMATION SERVICE

The goal of the core component Product Library is to provide the master product reference by consistent metadata and data management as well as data security. The data and metadata access is very flexible due to query based mechanism. For security reasons the access to the Product Library is well defined and restricted. To avoid attacks by internet users it is not possible to query the Product Library from public access points directly. The user information service is therefore strictly separated from the Product Library and the other DIMS core components. To provide selected metadata, browses and product data to the user there is a separate catalogue and data server in the user information service. The bridge between the DIMS core system and the catalogue and data server is the UIS-to-core interface (see Fig. 1). Via the UIS-to-core interface product metadata, browses and data are loaded to the user information service and are therefore accessible for public. The UIS-to-core interface also fetches user orders from the user information service and transfers them to the Order Management component.




Fig. 1 UIS-to-core interface bridges DIMS core system and user information service

Fig. 2 shows a more complex view of the user information service and the DIMS core components configuration. The persistence lines show data flows and the dashed lines show orders and requests. Ingestion and processing systems submit data at the Product Library. The Product Library triggers the UIS-to-core interfaces for uploading new products to the user information service and making them accessible by the user. The UIS-to-core interfaces also transfer orders from the user information service to the Order Management component, which sends processing requests to Production Control. Production control manages the workflows for processing and post-processing of the delivery products and finally assigns the Online/Offline Product Generation and Delivering component for generating the delivery units.

For offering products by a user information service there are several distinct requirements from near realtime availability to access restriction. Therefore the UIS-to-core interface has to provide several configurable mechanisms for uploading products. It is possible to upload products automatically by query based triggers or timers or to start upload actions manually. For both modes priorities can be specified. NOAA or MODIS products for example have to be available as soon as possible after acquisition. So the UIS-to-core interface is triggered by the Product Library as soon as a new acquisition is inserted in the catalogue and the data are archived. The trigger initiates the UIS-to-core interface to load the metadata and browses with high priority to the user information service. For other products which are less time critical like backlog processed products the upload can be configured with lower priority. The upload function can be restricted for the specific metadata and product components. An additional feature of the UIS-to-core interface is that image processing or data formatting can be applied to browses and product data during upload. A prominent use of this feature is the creation of thumbnails shown as a result of a user's product search. There are also specific configurations in the UIS-to-core interface component possible e.g. to automatically overwrite older products in the user information service whereas all versions are kept in the Product Library.




Fig. 2 Multi-Service Environment for providing Earth Observation Products to User Services

The UIS-to-core interface provides the possibility to upload products from one DIMS Product Library to multiple user information services or to upload from different DIMS Product Libraries to one user information service. For example GOME-2 Level 2 products (tracegas concentration) archived in the Product Library in Oberpfaffenhofen are loaded to the UMARF user interface of EUMETSAT. For MODIS the system is prepared to interface ESA multi-mission facility with DLR core systems. Uploads can be configured and monitored easily with the Operating Tool for the DIMS core services. After system downtimes the UIS-to-core interface checks automatically for new products in the Product Library and uploads them. The UIS-to-core interface has an own database for storing trigger and timer information as well as the upload queue with the upload status. Via the Monitoring and Alarm component of DIMS the operator is immediately informed when an upload query failed.

USER ACCESS TO EARTH OBSERVATION PRODUCTS

The user information service is the public access point to the data. Users want to have comfortable gateways. They want to be able to get quick information with visualisation for the products. They want to get all data from the same gateway and not use for each data collection a separate interface. So it's also necessary to upload products from different core systems to one user information service. For example in the user information service EOWEB there are metadata, browses and data offered from the Product Library in Oberpfaffenhofen and from the Product Library in Neustrelitz as well as data from old data archiving systems.

For retrieving products it is possible to get them directly via online transfer for a selected collection of small products or to order them for being delivered online via ftp or offline on media. The orders are submitted the user in the user information service and are stored in the user information service server. The UIS-to-core interface fetches the orders from the user interface service server and brings it to the core service Order Management. For security reasons the active part for the transfer of the order is as in the upload case the UIS-to-core interface which is located behind the firewall. Of course ordering is not only possible from the DFD user interface systems EOWEB, it is also possible from the EUMETSAT UMARF and from ESA MUIS.

More and more users are interested in higher level products as in low level ones. They want to get standardized products which they can use without further processing or formatting. It is not efficient to store each higher level product in each format in the archive if these products can easily be post-processed from the lower level products. So much products offered in the user interface system do not really exist, some offered collections are therefore post-processible on demand. For such products there are order options can be specified for delivery and their processing. Derived from the order options Order Management sends processing requests in the DIMS core subsystem Production Control to generate the ordered products. After the products are generated the subsystem Online/Offline Product Generation delivers them via FTP or generates the media (CD, DVD or DLT) for offline delivery.

EXPERIENCE

The German Remote Sensing Data Center is operating the DIMS core system together with the user information service EOWEB since 2000. The advantage to the former systems is that DIMS is multi mission capable. For all products of distinct missions inserted in the Product Library in Oberpfaffenhofen and offered via EOWEB the operator needs to monitor one Product Library and one UIS-to-core interface only. This is done with the unified Operating Tool which has the same look and feel for all components. By this way it was possible to reduce the operating effort significantly.

The UIS-to-core interface is able to upload huge amounts of data in the background. For the mission SRTM about 175 000 products (digital elevation models and interferometric data sets) have been uploaded systematically without problems. Due to the trigger mechanism of the UIS-to-core component MODIS products are accessible in the user information service about 10 minutes after they are ingested into the Product Library. The improvement to the former ingestion system is about

half an hour. The query based upload mechanism was a major improvement for example to satisfy project dependent access restriction.

FURTHER DEVELOPMENT

For the online transfer the products have to be stored in the user interface system server. To be able to delete older products and nevertheless make them accessible the next generation of the UIS-to-core systems gets the functionality for archive reloading into the online data cache of the user interface system. As a further feature it will be possible to upload configurations for product models and order options from the core services to the user interface system to reduce configuration and administration work.

APPLICATION OF NEW PRODUCTION STRUCTURES TO DECREASE ACCESS DELAY TO EO PRODUCTS

To improve the access time to user products a short view to the internal order of events in the DIMS has to be made. In DIMS the production control organizes processing chains to produce the ordered product. It activates the relevant processing systems and reports status back to production control.




Fig. 3 Conventional Data exchange between Processing Systems and Product Library

A so called processing system is the operational unit responsible for the processing of products of certain product types. It consists of a management part (processing system management - PSM) and the core processors. All administrative tasks like cache management and scheduling are performed by the PSM. It also provides the core processors with a needed temporary environment inclusive all needed input data. The actual processing is done by the core processors. In Fig. 3 this constellation is shown. Input data for a certain processing order are transferred from product library to the input partition of the core processor. Output products are handled in the opposite manner. In conventional production chains the input data provision and the output data take away is made by some kind of real transfer (e.g. FTP) in a LAN. Investigations [3] have shown that these real transfers have a considerable impact on total delay of production of EO data. In Fig. 4 this effect is shown for the L0 data production of CHAMP mission (from reception of a CHAMP dump up to the archiving and provision of the L0- product). Transfer and copy processes cause approximately a quarter of the total delay. Unbalanced investments in computing power for the processor can only decrease its 18%- share. Similar relations exist for a multitude of missions.

To weaken this effect SAN technology can be a useful way. Some well known advantages using a SAN are:

- Fast transfer speeds (The used Fibre Channel has a current transfer bandwidth of up to 400 MB/s by a physical transfer rate of up to 2 Gbits/s (current on the market))
- Improved scalability to expand the storage capability (up to 2^{24} N-ports possible)
- Interoperability between different systems
- Flexibility to connect and share remote systems
- Centralized management of the storage systems
- Systems with high reliability possible
- Cost-effective for large storage constellation

There exist a multitude of information about this subject and several solutions on the market. But not all applications use the full potential of this technology. First investigations in our institute were done last year. Best results were achieved using a shared file system constellation. Already in a not optimized test configuration the transfer performance was considerable increased (4x writing, 10x reading) against conventional Ethernet transfers.




Fig. 4 Sources of time consumption for the L0 processing of CHAMP [4] data [3]

The use of SAN technology has the potential for decreasing transmission delay not only by shortens the transmission itself but also by reducing the number of transmissions. If an archive based on hierarchical storage management (HSM) can share its disk cache with the processor environment several transmission can be skipped:

Input data provision	
Conventional:	Shared archiving and processing data partition
- Staging of needed data from archive media of the hierarchy into archive cache - FTP transfer to the processing system environment	- Staging of needed data from archive media of the hierarchy into archive cache - Mapping into the core processor environment
Output data take away	
- FTP transfer to the TMP area of the archive cache - Mapping of output data from TMP area into the archive controlled area - Archiving to media	- Output data directly written into the TMP area by the core processor - Mapping of output data from TMP area into the archive controlled area - Archiving to media

Fig. 5 shows this special use of a SAN. SANs also provides management tools to secure data access of archive and processing systems to avoid conflicts and data mutilation.




Fig. 5 New approach for data exchange between processing systems and product library

SUMMARY

The separation of user information service and core systems gives freedom to the specific functionalities of both domains. The UIS-to-core interface bridges these systems by providing features to exchange product and order information in a systematic and event driven way.

The increased integration of archive and EO data production via SAN can cause a considerable reduction of delays. This improvement comes along with new challenges to data security. But SAN technology also provides management tool to control the new structures.

References

- [1] Mikusch, E., Diedrich, E., Göhmann, M., Kiemle, S., Reck, C., Reißig, R., Schmidt, K., Wildegger, W., Wolfmüller, M.: Data Information and Management System for the Production, Archival and Distribution of Earth Observation Products. The DASIA 2000 Data Systems In Aerospace Conference, Montreal, 22.-26.05.2000, EUROSPACE, ESA, SP-457, pp. 401-406, 2000
- [2] Kiemle, S., Mikusch, E., Göhmann, M.: The Product Library – A Scalable Long-Term Storage Repository for Earth Observation Products. Data Systems in Aerospace 2001, EUROSPACE. ESA Publications Division, SP-483, Noordwijk (2001)
- [3] T. Hemmelskamp, „Optimierung des Datenmanagements am Nationalen Bodensegment Neustrelitz des Deutschen Fernerkundungsdatenzentrums“, Diplomarbeit, Technische Universität Braunschweig, September 2003
- [4] http://op.gfz-potsdam.de/champ/index_CHAMP.html

Integrating research data into the publication workflow: the eBank UK experience

Rachel Heery, Monica Duke, Michael Day, Liz Lyon⁽¹⁾, Michael B. Hursthouse, Jeremy G. Frey, Simon J. Coles⁽²⁾, Christopher Gutteridge, Leslie A. Carr⁽³⁾

⁽¹⁾UKOLN

*University of Bath, BA2 7AY, United Kingdom
Email: r.heery, m.duke, m.day, l.lyon@ukoln.ac.uk*

⁽²⁾School of Chemistry,

*University of Southampton, Southampton SO17 1BJ, United Kingdom
Email: m.b.hursthouse, j.g.frey, s.j.coles@soton.ac.uk*

⁽³⁾School of Electronics and Computer Science,

*University of Southampton, Southampton SO17 1BJ, United Kingdom
Email: c.jg, lac@ecs.soton.ac.uk*

INTRODUCTION

E-science has the potential to enrich scholarly communication from the perspective of both research and learning. This paper will explore emerging infrastructure that will enable effective curation of data to deliver this potential. The exponential growth in digital data arising from e-science requires new modes of data curation. The large volume of digital data now being created calls for new modes of publication, discovery, re-use and preservation. New patterns of **publication**, ‘publication at source’, are required, embedding the publication of data into the research scientist’s work patterns. Curation of digital data must support **discovery** services, building new service models on traditional library and publishing practice. There are opportunities for **re-use** of the increasing amount of digital data being produced, whether re-use takes place as part of research activity or in the development of course materials. In addition, secure archiving is necessary to ensure **preservation** of digital data.

The re-use and sharing of original data is a fundamental tenet of e-science. Successive levels of derived data will be produced from original data through refinement, transformation, interpretation, and generalisation. Original data is beginning to be made available in publicly available databases, whether institutional repositories, or publishers’ databases. This is happening particularly in certain specialisms, for example in the bio-informatics field, protein and genome sequences databases, and in chemistry, material safety information, and, significantly for eBank UK, crystal structures.

As new methods and technologies emerge to support the research data life-cycle through publication, discovery, re-use and preservation, there is a need to manage original and derived data (datasets produced both by experimentation and data re-use) alongside bibliographic data in an integrated fashion. Technologies underlying the emerging global network infrastructure (Internet, World Wide Web, Semantic Web, Grid) increasingly will support links between the activities of e-research and e-science, digital libraries and e-learning, offering enhancements to scholarly communications throughout the life-cycle from experimental research to learning. There is increasing interest in establishing institutional repositories to provide services to manage the variety of digital materials that form the intellectual output of educational and research institutions [1]. Significantly such institutional repositories also have a central role in taking forward the open access agenda, supporting e-print archives that enable self-archiving of journal article pre-prints and post-prints, and potentially other categories of material, by academic staff [2].

This changing landscape offers opportunities to curate scientific datasets more effectively, improving access and integration between services. The eBank UK project is addressing this challenge by investigating the role of aggregator services in linking metadata describing e-prints of peer reviewed journal articles to datasets from Grid-enabled projects made available within institutional and publisher e-data repositories.

THE EBANK UK PROJECT

The Joint Information Systems Committee (JISC) funded eBank UK project is led by UKOLN in partnership with the Universities of Southampton and Manchester. The project is working in the chemistry domain with the EPSRC (Engineering and Physical Sciences Research Council) funded e-science test-bed Combechem [3] at the University of Southampton, a pilot project that seeks to integrate existing structure and property data sources into an information and knowledge environment. The eBank UK demonstrator is being developed within a particular chemistry domain, crystallography, with a view to assessing, in the longer term, the feasibility of a generic approach across other disciplines.

The Combechem project offers an ideal research test-bed as it generates large quantities of digital data including crystallography data and physical chemistry data. Within chemistry research the vast majority of measurements in experiments are obtained using computer controlled instruments, so both data and metadata is provided automatically in digital form. Such data must be properly curated so a well-managed archive is vital, an archive that will outlast the completion of research projects and possible changes in staffing. Establishing institutional repositories for scientific data, maintained by a university or research institute, offers a possible solution. In order to populate such institutional e-data repositories effectively, the creation and archiving of both data and metadata must be fully integrated into the experimental workflow. Combechem is a test-bed for this approach, with a focus on crystallography and spectroscopic data as exemplars. eBank UK has chosen to focus on crystallography as this area of research has a strict workflow and produces data that is rigidly formatted to an internationally accepted standard, namely the Crystallographic Information File (CIF) standard [4], maintained by the International Union of Crystallography [5].

EBANK UK MOTIVATION

At present, as a result of technological advances in instrumentation and with the advent of e-science and high throughput philosophies, a publication bottleneck exists in many scientific communities. Accordingly only a small percentage of the data generated by many scientific experiments appears in, or is referenced by, the published literature. In addition, publication in the mainstream literature still offers only *indirect* (and often expensive) access to this data. As a consequence the user community is deprived of valuable information and funding bodies are getting a poor return for their investments. By moving towards a 'publication at source' approach access and discovery of research data would be made easier, which in turn would encourage re-use of data in further experiments.

Currently once a research experiment is finished the initial dissemination may be via a letter or communication, followed later by a more detailed explanation in a full paper. eBank UK will demonstrate how data might be 'published at source' directly in open archives and subsequently linked to either peer-reviewed journal articles or automatically deposited entries in specialised databases. Such an approach would enable the 'publication' of the enormous amount of data that is being generated within research areas such as crystallography. Publishers of crystal structures, and indeed most types of scientific data, make available to subscribers limited amounts of results data, such as CIFs. However this is only offering a partial solution, as the final result in most cases is merely a fraction of the digital data generated during the course of the experiment. Moreover this single result does not allow the 'reader' to assess the interpretations and assumptions made by the experimenter during the data workup. This is due to traditional publishing protocols being time consuming and considerably delaying, incapable of keeping up with the current data explosion. In addition such services are selective, and do not make available all the metadata that is associated with the primary data.

In parallel to increasing interest in archiving data, there is growing policy support and funding to encourage researchers to self-archive peer-reviewed journal articles in institutional repositories. This offers an opportunity to integrate services offered by e-data archives and journal article e-print archives. As the availability of journal article e-prints increases, users would benefit from a service linking back from the article to the raw or processed data. Various services might be built based on metadata made available through institutional repositories: providing context sensitive linking within electronic versions of journal, links from e-print archives to datasets, linking to datasets from e-learning portals.

The eBank UK demonstrator shows how an e-prints repository of peer-reviewed journal articles could include links to associated research data. In particular eBank UK focuses on the potential role for aggregator services that might harvest metadata describing datasets from institutional repositories as contributed by Grid-enabled projects.

EBANK UK ARCHITECTURE

The eBank UK project is exploring linking from metadata describing journal articles to the original datasets, thereby integrating access to datasets from Open Access e-print services. The aim of the project is to develop a pilot service.

This will be done by

- Creating an institutional repository of crystallography data, an open e-data repository (at the University of Southampton)
- Modifying e-prints.org software [6] both to support the open e-data archive, and to enable harvesting of metadata describing datasets (software developed and maintained by the University of Southampton)
- Demonstrating an eBank UK search service linked to ePrints UK [7], indexing harvested descriptions of datasets and Journal articles (at UKOLN, University of Bath)
- Embedding the eBank UK service into the PSIGate [8] subject gateway (at the University of Manchester)

The architecture of eBank UK adopts the Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH) [9] which offers a specific design approach. OAI-PMH envisages a class of ‘data providers’ acting as repositories of resources. These data providers are accessed programmatically by means of a number of requests (which together make up the OAI-PMH). A second class of ‘service providers’ issues the OAI-PMH requests to the data providers. A series of requests made by the service provider (the full details of which are beyond the scope of this paper) culminates in ‘harvesting’ of the metadata about the resources held by the data provider. In other words, resource descriptions held by the data provider are sent in response to the service provider requests. By selectively making requests to one or more data providers, the service provider aggregates a collection of resource descriptions. The service provider (or aggregator) is thus positioned to provide a single point of entry to disparate repositories of resources and resource descriptions, and can offer a variety of services to support the discovery of published data and literature.

The harvesting design addresses some well-known limitations of cross-searching techniques. Cross-searching techniques generally send specific search requests in parallel to different sources (by some specified protocol) and combine the various responses into a result for the cross-search. In contrast, search services built on harvested metadata carry out local searches on the pre-harvested metadata. This circumvents two problems which may be encountered in cross-searching: firstly delayed responses necessitate that the requester decides whether to wait on all responses before presenting any results (for example before applying any ranking) and secondly the potential unavailability of the search target at the moment the request is made may leave a gap in the response, corresponding to the unavailable target.

For historical reasons, the OAI-PMH is very closely tied with the Open Access philosophy of publication, whereby the dissemination of publications (such as journal articles) is improved by depositing the item in an institutional or subject-based repository. The repository acts as a central focus from which the article, together with the associated metadata, may be shared freely with the research community. The very first Open Access collections (e.g., arXiv [10], CogPrints [11]) provided the test-bed on which the OAI-PMH was developed and refined, until it evolved into the current stable version. Partly as a consequence of this association, OAI-PMH-compliant repositories have proliferated, and a number of Open Access collections fulfilling the role of data providers (in an OAI-PMH sense) are currently available. These data providers are distributed in the US, Europe, Australia and wider. They take on a different focus according to their context; many are institutional and are seen as a central point of dissemination for an institution’s published assets; however different flavours of repositories are possible, for example centred on images, theses and even CVs.

Compared to the number of data providers, the emergence of service providers has been slower. Arguably this is an unavoidable situation. A necessary level of data provider availability is required to provide the material to fulfil service provider requests. A range of data provider instances would also be needed to motivate the development of different service provider models. A lag-time may be inevitable and service provider maturity may follow on the success of the data provider take-up.

As stated, service providers act as ‘aggregators’, selectively harvesting metadata from one or more data repositories. The OAI-PMH enables a degree of selective harvesting based on a number of factors. At the most basic (and mandatory) level, OAI-PMH supports requests based on the last-updated date, so that cumulative harvesting of metadata can be performed. Resources can also be optionally partitioned into ‘sets’, defined according to a manner appropriate to the data provider (e.g. by subject or by resource type). Finally, service providers can be selective in deciding which data providers to harvest, for example only harvesting from collections that meet quality criteria by enforcing policies about the material deposited. Alternatively a service provider could concentrate on meeting the

needs of a specific community with an interest in certain kinds of resources. Furthermore, the service provider can tailor the searching services to the searching and discovery requirements of the particular user group that it serves. In this manner, service providers can differentiate the services on offer.

The ePrints UK service, funded by JISC, is providing a focus for the cross-searching of assets made available by institutions, presenting a subject-based service targeting the existing user base of JISC-funded subject-based resource discovery services. Furthermore, e-prints UK is seeking to enhance the harvested metadata through citation analysis, name authority services and automated subject classification.

In the eBank UK project, this Open Access philosophy is being extended to research data. A data repository has been created at the University of Southampton. Crystallographers deposit the datasets produced at the National Crystallography Service in an enhanced version of the e-prints.org software that accommodates self-archiving by chemists. During the deposit process, metadata about the datasets is entered or generated automatically. This metadata is then made available in the repository that, in addition to supporting local searching, browsing and downloading of datasets, also acts as an OAI-PMH compliant data provider. eBank UK has also implemented an example of a service provider that harvests metadata about datasets from the repository (using OAI-PMH). This metadata provides the basis on which to create innovative services that support the discovery and re-use of datasets, as well as creating links with the published literature. The shared infrastructure re-using the OAI-PMH supports interoperability between systems disseminating publications and those disseminating research datasets.

EBANK UK DATA FLOW

Fig. 1 summarises the flow of data and metadata within the eBank UK project. Datasets are created during experimental processes and are instantiated as a number of data files. The files are self-archived in an OAI-PMH repository and descriptive metadata is created as part of the deposit process. The metadata consists of a number of fields that are either entered either by the depositor or generated automatically by the modified e-prints.org software. In the institutional eData repository, a local web interface (using HTML) is presented to the user. This interface consists of crystallographic eData reports, an interactive visualization of a derived crystal structure, selected values that describe the experiment and the results, as well as links to all the data generated in the course of an experiment.

The eBank UK aggregator harvests the metadata from this data repository by making OAI-PMH requests. The metadata schema used for exchange is discussed in the next section. An SGML indexing and searching engine, Cheshire [12], is then used to provide indexing and searching functionality over the metadata. When combined with




Fig. 1. Data flow in the eBank UK architecture

publication metadata, any links made between different datasets, or datasets and publication can be revealed to the user in the display of the search results. Links can be made indirectly, for example by identifying common occurrences of author names, keywords or other subject-specific terms (such as chemical identifiers in our particular crystallography exemplar). Alternatively, direct links can be created if there is a specific reference, in the dataset or the publication metadata, to the related work. For example, an OpenURL could be generated to link to an instance of a journal article that is available electronically in a user's institution, if sufficient metadata about this related article is included in the metadata record describing the data set.

One further area of work in the eBank UK project is the embedding of the search service into external services to reach the user base by means of alternative points of entry. There are a number of mechanisms by which this can be achieved. All the mechanisms would allow a third party service to make requests to the service provider, which in turn returns results in a manner that can be manipulated by the third party to fit the look and feel of their interface, for example the University portal or a subject portal. EBank UK has implementation experience of using CGI-mechanisms, web service protocols (SRW Search/Retrieve Web Service [13]) and portal technologies. These techniques will be used to present the search service through PSIGate, which has an audience in higher and further education, mainly in the UK, but also worldwide, with an interest in the physical science. The eBank UK search service will complement the other features of PSIGate, adding discovery and access of research data to the landscape of scientific services already on offer at PSIGate.

EBANK SCHEMAS

The services that can be built on aggregated metadata can only be as good as the metadata that is available to them. To achieve interoperability between dataset and publication metadata, there must exist some consensus on the data models and the metadata schemas being used to exchange metadata. Furthermore, to support discovery facilities between datasets from different communities, there must be some agreed commonality in the models and schemas if cross-disciplinary services are to be built.




Fig. 2 . The links between metadata records, HTML views of the records and data sets (left) and the eBank UK architecture (right)

As a starting point the eBank UK project has taken the Dublin Core metadata element set [14], and other DC terms available [15], and adapted them to describe the datasets that are deposited by crystallographers in the institutional eData repository. The reasons for this are two-fold. Firstly, simple DC is mandated within the OAI-PMH (although the protocol also encourages the exposure of specialized metadata records alongside the minimally required DC). Secondly, DC was conceived to provide a minimum baseline of interoperability, which can be extended and specialised for specific needs.

The metadata schema uses elements from the fifteen elements of the Dublin Core Element Set as the basis for descriptions. For example the creators of the datasets are designated within the Dublin Core creator element. Discussion with the users revealed that there are a number of different ways to describe the subject of the datasets. Crystallography experiments revolve around a single molecule, which can be thought of as the 'topic' of the experiments. There are a number of established ways of identifying molecules, which include internationally recognised methods of specifying their formulas or names. These different vocabularies have been incorporated into the schema through the encoding schemes facility of qualified Dublin Core. The Dublin Core Metadata Initiative (DCMI) provides recommendations for including terms from vocabularies in encoded XML suggesting "*Encoding schemes should be implemented using the 'xsi:type' attribute of the XML element for the property.*" An example from the eBank UK metadata records for datasets is the representation of the chemical empirical formula as:

```
<dc:subject xsi:type="eBank UKterms:empiricalFormula">C288 H200 Cl24 F48 O48 P16 Pd16</dc:subject>
```

Efforts are ongoing in the chemistry community to agree on namespaces for these vocabularies [16]. As these namespaces emerge, they can substitute the eBank UK namespace designations, which have been used temporarily until the standard recommendations become available.

The metadata record harvested by the eBank UK aggregator describes the collection of datasets which are identified by a URL that links to the HTML entry page in the data repository where they were deposited. This identifier is used as the value of the DC identifier element. The location for individual datasets can also be given in DC relation elements. Since it is desirable to be able to describe the type of datasets, the extended DC vocabulary mechanism was once again used to specify types.

```
<dc:type xsi:type="eBank UKterms:EBank UKDatasetType">CollectionDataset</dc:type>
```

Related literature, such as journal articles, can also be exposed in the metadata record, either by citing a unique identifier (e.g. DOI) or by textual citation (author details, title, issue etc.).

However this metadata schema (and the Dublin Core model) does not accommodate the whole complexity of the model of datasets. There may be some inherent relationships between datasets that would be usefully revealed in the metadata. For example in the case of crystallography, datasets are related to one another by sequence since they are generated (by




Fig. 3. The crystallography work flow and relationship between data sets

measurement or analysis) from a series of consequent stages in the experimental process. Similarly, datasets may be made up of one or more files, which may again be related in specific ways, or crucially, may be stored in different locations or facilities, governed by varying access control mechanisms.

The CCLRC data model [17] developed by the Council for the Central Laboratory of the Research Councils, is an emergent example that attempts to deal with this complexity by describing the relationship between experiments, investigators, data holdings, datasets, data files, logical and physical locations. Another area of interest is provided by recently proposed schemas for describing complex objects, such as the Metadata Encoding and Transmission Standard (METS), the MPEG21 Digital Item Declaration Language (DIDL) [18] and content packaging standards from the e-learning community.

CONCLUSION

There are increasing commonalities in the concerns of those developing digital libraries and those creating and managing e-science data. This includes the potential use of common technologies for managing datasets and bibliographic data e.g. the OAI-PMH; and overlaps between services that might be offered based on the data e.g. linking datasets and journal articles

During its initial phase, eBank UK has focused exclusively on the chemistry domain and in particular on crystallography. There is potential to expand its remit to a wider range of crystallography data, to other chemistry sub-domains and indeed to encompass the broader range of physical sciences. A longer term ambition would be for other scientific communities, including social sciences and humanities, to adopt a similar approach. eData repositories might be supported, both by institutions (universities and research institutes) and publishers. In order to progress this vision an immediate challenge for the wider community is to reach consensus on a common data model for scientific datasets.

Agreement would be required on a common approach to the range of services that might be built on such repositories. Value added services might include the enhancement of data with visualisation and scientific context through the use of mark up languages (e.g., Chemical Markup Language (CML) [19, 20], Computational Chemistry Markup Language (CCML), Mathematical Markup Language (MathML) [21]). Other services more typically associated with digital libraries might be offered, such as enhancement of metadata by automated subject classification and authority control. The provision of services connected with subject access such as these, would be particularly appropriate in the context of subject portals, and eBank UK is indeed exploring this area in collaboration with the Resource Discovery Network PSIgate service.

There are other outcomes resulting from enhancements to curation of scientific data that are outside the scope of this paper. These include the pedagogical benefits of providing access to primary research data within eLearning materials. Within the context of eBank UK, this would be possible within postgraduate courses in chemistry, undergraduate project work, or chemical informatics courses. In a wider context there could be inclusion of e-research data in e-learning courses, through links in reading lists, essay assignments, analytical problem solving, and through practical work. Enhanced availability of and access to original and derived scientific data might also suggest possible changes in the peer review process. The provision of all associated data in a repository might satisfy most scientists that the formal refereeing of data via peer review could be removed from the publication chain, since the data can be assessed readily on line by anyone at any time.

In order to move from demonstrator to service, there is a need to involve publishers and other existing specialist services, as well as institutions, in the eBank UK approach. There is potential for a variety of business models whereby publishers might build services based on harvested metadata, particularly those publishers that already have a focus on access to original data.

In conclusion, the eBank UK project has built on a joint approach arising from both the digital library community, the Grid and computer science community. eBank UK demonstrates benefits to the research community, and has shown the potential for integration into digital library services. For more information on the eBank UK project, please see the project web pages at <http://www.ukoln.ac.uk/projects/ebank-uk>.

ACKNOWLEDGEMENT

The eBank UK project is funded by JISC under the Semantic Grid and Autonomic Computing Programme.

REFERENCES

- [1] C. Lynch, "Institutional repositories: essential infrastructure for scholarship in the digital age," *ARL Bimonthly Report*, no. 226, February 2003. Retrieved July 22, 2004, from: <http://www.arl.org/newsltr/226/ir.html>
- [2] R. Crow, *The case for institutional repositories: a SPARC position paper*. Washington, D.C.: Scholarly Publishing & Academic Resources Coalition, 2002. Retrieved July 22, 2004, from: <http://www.arl.org/sparc/IR/ir.html>
- [3] J.G. Frey, M. Bradley, J.W. Essex, M.B. Hursthouse, S.M. Lewis, M.M. Luck, *et al.*, "Combinatorial chemistry and the Grid," in *Grid computing: making the global infrastructure a reality*. F. Berman, G. Fox and A.J.G. Hey, Eds. Chichester: Wiley, pp. 945-962, 2003.
- [4] S.R. Hall, F.H. Allen and I.D. Brown, "The Crystallographic Information File: a new standard archive file for crystallography," *Acta Cryst.*, vol. A47, pp. 655-685, 1991.
- [5] International Union of Crystallography. Retrieved July 22, 2004, from: <http://www.iucr.org/>
- [6] The eprints.org software. Retrieved July 22, 2004, from: <http://software.eprints.org/>
- [7] ePrints UK project. Retrieved July 22, 2004, from: <http://www.rdn.ac.uk/projects/eprints-uk/>
- [8] The PSIGate service Retrieved July 22, 2004, from <http://www.psigate.ac.uk/>
- [9] C. Lagoze, H. Van de Sompel, M. Nelson and S. Warner, Eds., *Open Archives Initiative Protocol for Metadata Harvesting*, version 2.0, June 2002. Retrieved July 22, 2004, from: <http://www.openarchives.org/>
- [10] arXiv.org e-Print archive. Retrieved July 22, 2004, from: <http://arxiv.org/>
- [11] CogPrints. Retrieved July 22, 2004, from: <http://cogprints.ecs.soton.ac.uk/>
- [12] Cheshire SGML search engine. Retrieved July 22, 2004, from: <http://cheshire.berkeley.edu/>
- [13] SRW Search/Retrieve Web Service. Retrieved July 22, 2004, from: <http://www.loc.gov/z3950/agency/zing/srw/>
- [14] Dublin Core Metadata Initiative, *Dublin Core Metadata Element Set*, version 1.1, DCMI Recommendation, June 2003. Retrieved July 22, 2004, from: <http://dublincore.org/documents/dces/>
- [15] Dublin Core Metadata Initiative, *DCMI Metadata Terms*, DCMI Recommendation, June 2004. Retrieved July 22, 2004, from: <http://dublincore.org/documents/dcmi-terms/>
- [16] XML Data Dictionaries in Chemistry. Retrieved July 22, 2004, from <http://www.iupac.org/projects/2002/2002-022-1-024.html>
- [17] S. Sufi, B. Matthews and K. Kleese van Dam, "An interdisciplinary model for the representation of scientific studies and associated data holdings," UK e-Science All Hands Meeting, Nottingham, UK, 2-4 September 2003. Retrieved July 22, 2004, from: <http://www.nesc.ac.uk/events/ahm2003/AHMCD/pdf/020.pdf>
- [18] J. Bekaert, P. Hochstenbach and H. Van de Sompel, "Using MPEG-21 DIDL to represent complex digital objects in the Los Alamos National Library Digital Library," *D-Lib Magazine*, vol. 9, no. 11, November 2003. Retrieved July 22, 2004, from: <http://www.dlib.org/dlib/november03/bekaert/11bekaert.html>
- [19] P. Murray-Rust and H.S. Rzepa, "Chemical markup, XML, and the Worldwide Web, 1. basic principles," *J. Chem. Inf. Comput. Sci.*, vol. 39, pp. 928-942, 1999.
- [20] P. Murray-Rust and H.S. Rzepa, "Chemical markup, XML, and the World Wide Web, 4. CML Schema," *J. Chem. Inf. Comput. Sci.*, vol. 43, pp. 757-772, 2003.
- [21] D. Carlisle, P. Ion, R. Miner and N. Poppelier, Eds., *Mathematical Markup Language (MathML) Version 2.0*, 2nd ed. W3C Recommendation, October 2003. Retrieved July 22, 2004, from: <http://www.w3.org/TR/MathML2/>

Added-value services in Ether atmospheric chemistry data centre

Cathy BOONNE ⁽¹⁾, Françoise GIROD ⁽²⁾

⁽¹⁾ IPSL, Institut Pierre Simon Laplace
Université Pierre et Marie Curie
T46, 4, place Jussieu
75252 Paris
Email: Cathy.Boonne@ipsl.jussieu.fr

⁽²⁾ CNES
DCT/PS/VDO
18, avenue Edouard Belin
31401 Toulouse Cedex4
Email: Francoise.Girod@cnes.fr

INTRODUCTION

The atmospheric chemistry experiments provide a huge amount of data available from many different data centre. The data files content analyze often shows that the same chemistry variable is expressed in various units and references from one experiment to another. The consequence of the lack of common standard as regards units or references is that the use of some experiment data is no easy to achieve. In Ether, the French atmospheric chemistry data centre, the same problems are encountered when scientists want to compare, validate or visualize data from many sources. Data normalization is an answer to solve these recurrent problems.

ETHER ATMOSPHERIC CHEMISTRY DATA CENTRE OVERVIEW

The French atmospheric chemistry data centre Ether has been developed and funded by the French Space Agency (CNES) and the Institut National des Sciences de l'Univers (INSU/CNRS) for about 4 years. The role of Ether [1] is to assist French atmospheric researchers and European scientists, as long as they were involved in co-operation agreements with French scientists, to locate, access and interpret atmospheric data.

This centre gathers data from satellite, balloon campaigns linked to satellite validation, aircraft and model results for stratospheric and tropospheric purposes and from different levels of production (raw data, physical data, interpolated or assimilated data).

At the beginning of 2004, Ether manages two sets of data:

- Public products : POAM II, SAGE II, UARS/WINDII, all chemistry UARS data (MLS, CLAES, ISAMS, HALOE), ADEOS1/ILAS and ADEOS/IMG,
- Restricted products: POAM III, ODIN/SMR, ODIN/OSIRIS, ENVISAT/GOMOS. These data products should be public in the next years according to the science team data protocols.

Ether contains also data provided by ground-based stations, balloon gondolas or aircraft measurement instruments (SAOZ, ELHYSA, LPMA-DOAS, SALOMON, and SPIRALE) and ancillary products.

Ether is managed by the IPSL organisation (a CNRS-INSU/IPSL institute), all data are available on-line through a World Web interface (<http://ether.ipsl.jussieu.fr>), see Fig. 1.

Software and added-value services (quick look, data format conversion, retrieval software, atmospheric models) are provided to assist in the manipulation of the data or to generate higher levels standard data products. Extensive information is also provided on the data collection procedures, formats, contact names and references to journal papers.

A group of scientific expert decides the Ether's evolutions as regards the handling and development of added-value services. These software must be of large interest for all Ether users, like Arletty software for simulation of atmospheric temperature and pressure or Msdol for the assimilation of ENVISAT/GOMOS and ODIN/SMR products which has generated assimilated data (level 4) 3 months after the end of the commissioning phase of the sensors.

Finally, the great advantage for scientific users of using Ether is also to have the capabilities to run on-line software and added-value services through a World Web interface; e.g. the processing software Moliere permits the inversion of Odin Level 1 data and the production of Odin Level 2 data files.

In addition to data support activities, a major goal of Ether is to promote the creation of different expertise networks on varying atmospheric chemistry topic. The first one is on data assimilation. Emphasis is placed on user involvement with Ether.



Fig.1. The Ether Atmospheric Chemistry data centre home page.

DATA NORMALIZATION [3]

Ether Context

Data Formats Managed By Ether Database

Data available through Ether database are collected from different and various data suppliers (scientific organisms, research laboratories, data provider centre, scientists ...) using different ways: ftp site, web site, mail, magnetic tape, etc. Atmospheric chemistry data formats are various and depend on the description defined by the principal investigator and the scientific project group of the relevant scientific experiment.

Ether database doesn't require a particular data format for the handling and management of data. Much more than thirty different data formats in Ether have been counted.

Diversity especially in level 2 products formats is encountered; for example in:

- Product type : we can find total column or profile,
- Product unit : it depends on the definition given by the scientific project group
 - ozone profile expressed in molecule.cm-3 or Dobson unit or ppv or ppmv ...
 - pressure profile expressed in atm or hPa or mb ...
 - temperature profile expressed in °C or K
 -
- Product naming: no product-naming convention, e.g. O3_B9, O3_5018 ...
- Latitude / longitude convention different from one experiment to another,
- Vertical level expressed in altitude or pressure,
- Data files format: ASCII, binary, HDF, AMES ...
- Grid projection,
- Distribution of the products information in one or several files.

The lesson learned about atmospheric chemistry data formats is that there are no real standard data formats clearly defined by the international atmospheric community.

Standard Data Format Requirements

In case of data validation, analyses, or visualization using data from many sources, scientists need to access rapidly to data which correspond to the date, time, latitude and longitude, same products units and geographical references. To solve this problem, the first idea was to define and develop a common format and uniform units for all the data sources. A standard format named “Ether” [2] was described and tested in 1999 solving normally all the problems encountered with the use of the native data format. Data products for four UARS experiences have been converted but the conversion effort has taken too much time (a few months for one version) for a real and significant scientific use. An other drawback of using this method is the necessity of re-generating standard files at each new data version update. Besides it appears a strong consensus among the French scientific community to use exclusively native products format to achieve their scientific work and to have facilities to exchange data with others data co-investigators. Scientists also promote consequently the development of conversion tools to harmonize the different data formats. The standard data format “Ether” was abandoned but not the need of data normalization in Ether.

Towards Data Normalization

Ether Proposal

The first step towards data normalization is to choose the relevant information in data files to standardize. In Ether database, a metadata catalogue describes the product information of all different datasets. Eight fields belonging to this metadata catalogue are chosen to be normalized; they are listed below in Table 1. For each field chosen, the idea is to define a standard format.

Attributes	Standard format
Date	YYYY/MM/JJ:HH:MM:SSZ
latitude	Degrees [-90° +90°]
longitude	Degrees [-180° +180°]
altitude	km
CEOS data level	L0, L1, L2, ...
Format	AMES, HDF, BINARY, ASCII, ...
Version	Version of the native data
Chemistry parameters	HITRAN convention

Table 1. Normalized attributes used in Ether database.

For chemistry parameters, it is possible to add a non-normalized sub-parameter to distinguish the selected parameter among several datasets. The data version convention is the same as defined by the data suppliers. For each datasets, it is also possible to add additional parameters, the definition of these depending on the product type or on the scientific use (orbit number, flight number, station name ...).

The normalization of these attributes is fundamental to fulfil the Ether users' requirements in data analyses, validation or visualization. Using it, it is possible to develop conversion tools to harmonize the data format of all Ether data sources.

Steps To Elaborate Conversion Tools

Step 1: Identification of the data files information

In the context of Ether, the first step to elaborate conversion tools using the standardization information defined in Table 1 is to identify all the Ether data files information, product types and units. Table 2 gives an idea of the ISAMS data files information.

Experiment	Contents of data file	Product type or format	Unit	File format
ISAMS	Date	JJ-'MMM'-YYYY		SFDU and ASCII
	hours	hh:mm:ss		
	latitude		degrees	
	longitude		degrees	
	solar zenithal angle		degrees	
	aerosol extinction coefficient or	profile	km ⁻¹	
	volume mixing ratio of one	10e-6 ppmv	vmr	
	chemistry parameter			
	Temperature		K	

Table 2: ISAMS data files information.

Step 2: Creation of a dataset descriptor and conversion tools

A descriptor of each datasets is built in a second step according to the schema in Fig.2.




Fig. 2. Ether data descriptor

The descriptor consists of the normalized attributes defined in Table 1 and of the additional data files information given in Table 2. For each datasets, the descriptor allows to combine chemistry parameter with the related standard format, and permits to define missing values. Using the descriptor, it is easy to develop conversion tools integrated in a or several data processing chains to generate standardized products or new products.

According to the scientific requirements, conversion tools are developed for:

- reading data files,
- calculating missing values if desired,
- converting data from one unit to another,
- changing of geographical references,
-

Step 3: Development of data processing chains

In a third step, data processing chains can be developed for the use of scientific work. These data processing chains can be used and associated in each new added-value services developed in Ether.

An example of a data processing chain is given in Fig.3.




Fig.3. Example of a data processing chain.

Installation and availability

Currently data normalization in Ether is on going. The different conversion tools and a data processing chain for visualization are under development. A visualization test interface using normalization tools will be installed on Ether system and will be available to Ether users on next September. Ether users will have the choice to download the tools to work locally or to use the tools developed through the World Web Ether interface.

CONCLUSION

Normalization of Ether data sources is fundamental for scientists in case of data validation, analyses, or visualization. Using data normalization, scientists gain time to access the data and to manipulate them using the same standard. Many data processing chains will be developed and available through Ether to help scientists to use efficiently data from many sources. In addition to that a strong effort will be shortly made in Ether to study the normalization of assimilated atmospheric chemistry data fields in order to define a level 4 standard format.

REFERENCES

- [1] P. Ricaud, C. Boone, F. Girod and T. Phulpin, 'Le pôle thématique de chimie atmosphérique Ether', 7 p, 2003.
- [2] F. Girod, 'Spécification des fichiers gérés et produits par Ether', 22 p, 1999.
- [3] F. Girod and C. Boone, 'Outils de normalisation des données. Spécifications techniques et conditions de réalisation.' 15 p, 2003.

Information Mining and Understanding of Data Massive: a new era of methodology and technology

Mihai Datcu, German Aerospace Center DLR Oberpfafenhofen
mihai.datcu@dlr.de

Klaus Seidel, Swiss Federal Institute of technology ETH Zurich
seidel@vision.ee.ethz.ch

Sergio D'Elia, ESA/ESRIN Frascati
Sergio.DElia@esa.int

Abstract: The decreasing cost of fast computers and large storage devices and the increasing speed of networks enable the management of large amounts of EO data, information and knowledge. However the capacity to understand data is further rather limited. Thus, an important and urgent priority is to develop new methodologies to metamorphose the computer from a machine used for solving difficult problems, into effective means for communicating with other people and machines. The concept to implement such a system, is to split the task in the most appropriate way to exploit the computer skills and to adapt to the human behavior.

The article presents new concepts and method to structure very large volumes of heterogeneous data, like images, multidimensional signals, text, or other types of records. Presents the main characteristics of processing very large volumes of data, and analysis:

- the technological complexity
- the computational load of solving typical tasks
- the characterization of the information content, and its complexity
- the limits of learning from small data samples
- the specific human perception aspects

The article concludes with the discussion of the perspectives of the information processing and data understanding in the fields of EO in relation with other branches of science and engineering.

Semantic Web Technologies for Searching, Retrieving and Adding Value to Large Scale Scientific Data

Catherine Houstis, Marios Pitikakis, George V. Vasilakis and Manolis Vavalis

Affiliation

Informatics and Telematics Institute, Centre for Research and Technology Hellas,

1st Km Thermi-Panorama road, 57001 Thermi, Thessaloniki, Greece

Email: {houstis,pitikak,vasilak,mav}@iti.gr

ABSTRACT

In this paper we discuss the relationship between the Semantic Web and the grid, and explain how grid applications can also be applications of the Semantic Web; this is illustrated by an environmental application scenario. We indicate how Semantic Web technologies can be applied to grid computing and we outline this approach by presenting an e-Science middleware system, ARION, that has been designed and implemented to support search and retrieval of scientific information.

INTRODUCTION

The scientific community has embraced the Web, enabling researchers as well as industrial communities to closely collaborate and share resources. The result is typically the publication of information without any accompanying accessibility. Scientific knowledge is often embodied in the literature and in free text “annotations” attached to raw data. The presumption that a scientist will read and interpret the texts makes automatic processing hard and is not sustainable given the increasing amount of data becoming available.

The Semantic Web is about sharing knowledge, making the computationally inaccessible accessible and about simplifying information discovery and making it more intuitive. Even simple metadata and simple queries give a small but significant improvement in information integration. Automated processing of Web content requires explicit machine-processable semantics associated with Web resources to describe what it is about and what it is for. The ambition interwoven with the Semantic Web is of an environment where software agents are able to dynamically discover, interrogate and interoperate with resources, building and disbanding virtual problem solving environments, discovering new facts, and performing sophisticated tasks on behalf of humans [9]. On the other hand, the ultimate goal of the Grid [3,4,5] is the pooling and coordinated use of large sets of distributed resources. There is a close relationship between these two technologies and each one can greatly benefit from the other.

In this paper, we will outline a middleware architecture that combines concepts found in both the Semantic Web and Grid visions. In our approach towards e-Science, we have already implemented a prototype system, ARION [11,12]. The main motivation for our work is the need to enhance *information-sensitive* systems with better mechanisms for e-science composition and integration, for data and resource management, and for real time requirements and integration of actions of different organizations (like ones for risk/crisis management). ARION is a service-based infrastructure designed to support search and retrieval of scientific objects, and capable of integrating collections of datasets and scientific applications including simulation models and associated tools for statistical analysis and dataset visualization. It also actively supports on-demand scientific data processing workflows, in both interactive and batch mode. The system’s demonstration scenarios involve mainly environmental applications (offshore to near shore transformation of wave conditions, synthetic time series and monthly statistical parameters, coupled ocean-atmosphere models etc.).

THE SEMANTIC WEB AND THE GRID

Until very recently the Grid and the Semantic Web communities were separate, despite the convergence of their respective visions. Both have a need for computationally accessible and sharable metadata to support automated information discovery, integration and aggregation. Both operate in global, distributed and changeable environments. The Semantic

Web could make its way into Grid core Services. The Semantic Web fabric is the means by which the Grid could represent metadata and semantics: both for Grid *infrastructure*, driving the machinery of the Grid fabric, and its core and high level services, and for Grid *applications*, representing the knowledge and operational know-how of the application domain.

The Grid Computing technology is the natural evolution of distributed computing with more emphasis on open architecture models, and simpler forms of integration of resources. The Grid Computing evolution will enable easier resource discovery and management, virtualization, and enhanced functionality for distributed applications.

The Semantic Grid [16] is a natural evolution of Grid Computing, towards a knowledge-centric and metadata-driven computing paradigm. The Semantic Grid is an effort to utilize Semantic Web technologies within Grid Computing development efforts, from grid infrastructure core services to the delivery of grid applications.

These concepts will further enhance automated resource and knowledge/information-based resource sharing. Realizing the importance of the evolution of the Semantic Web and its potential role in Grid Computing, the Global Grid Forum [6] has created a research group for Semantic Grid under the Grid Architecture area.

Grid and Semantic Web technologies appear symbiotic and their visions are related. Grid computing benefits from the Semantic Web for the management of semantics. The Semantic Web benefits from the application driving force provided by the Grid and the Grid infrastructure itself. The base services of the Semantic Web – ontology servers, metadata generators and so on – could be implemented as Grid Services.

A SERVICE ORIENTED VIEW

A great number of adding value services can be provided based on data processing of raw scientific data. Such services range from simple statistical analysis of data (low level processing) to decision support, expert system services, environmental policy scenarios etc. Visualization can be regarded as part of the data processing added value chain.

A service can simply be viewed as an abstract characterization and encapsulation of some content or processing capabilities. Thus, services can be related to the data, computational or knowledge grid. The data grid provides access (search and retrieval) to datasets stored in geographically distributed systems in various organizations. The computational grid deals with the way that computational resources are allocated scheduled and executed and the way in which data is shipped between the various processing resources. The knowledge grid handles the way that information (i.e. metadata, ontology) is represented, stored, shared and maintained. Given its key role in many scientific activities, the Web is the obvious point of interest at this level. Here information is understood as data equipped with meaning.

In designing such a system, a service oriented model must be adopted, supporting a high degree of automation and providing a flexible environment for large scale computations and resource retrieval. In a computing paradigm, a user might proceed by

- **gaining** the necessary authentication credentials and access rights
- **querying** the information system to determine the availability of required datasets and view the corresponding metadata information
- **submitting** requests to initiate computations (workflows), move data, and so forth
- **monitoring** the progress of the various computations, notifying the user when all are completed, and detecting and responding to failure conditions
- **visualizing** the results of computational workflows

All of the above give rise to the view of the e-Science infrastructure as a set of services that are provided by particular individuals or institutions for consumption by others, under various forms of contract.

THE ARION PROTOTYPE

ARION provides the means for organizing geographically distributed and heterogeneous resources of scientific content, so that its disparate and varied parts are integrated into a coherent whole. Our architecture combines both a horizontal and a vertical integration. The horizontal integration is expressed by ontologies and the vertical by workflows (applications).

These are the main mechanisms for recording expert knowledge, for information representation and navigation, and for expressing computation processes over a pool of resources.

ARION is composed of a set of distributed nodes containing data sets and programs (scientific collections). These nodes interoperate using an agent platform, and provide the basic infrastructure for workflow execution. Workflows rely on distributed and autonomous tasks and are controlled by a centralized server. Mobile agents installed on each node execute workflow tasks (mainly computations) and monitor the execution flow.

The architecture of our prototype consists of four subsystems. The *Search Engine* allows users to pose queries to the knowledge provided by ARION. It is mainly based upon RDFSuite [1], a suite of tools for RDF metadata management providing storage and querying for both RDF descriptions and schemas. It consists of three main components: a RDF validating parser (VRP), a RDF schema-specific storage database (RSSDB) and a query language (RQL). The RDF Query Language (RQL) [13] is used to uniformly query RDF metadata information and RDF schemas (i.e. queries can be posed for either the metadata information or the concepts and properties of the ontology). Thus, we can exploit this ability to implement schema browsing, since large RDF schemas carry valuable information themselves due to class refinement.

The *Workflow Data Base System* contains the workflow specifications and handles the preparation of execution specifications to be sent to the workflow runtime. It consists of the following components: Workflow Editor, Workflow Storage System, Workflow Database Server and Statistical Database.

The *Workflow Runtime System* is responsible for the execution of workflows and the management of the information produced during this execution from the distributed nodes. It consists of two main components, namely the Workflow Manager and the User Monitoring System. For each workflow definition received from the *Workflow Data Base*, an execution environment is initialized and a Task Scheduler is created. The Task Scheduler makes decisions about the order of execution of tasks and assigns blocks of tasks to the Task Manager, which is responsible for the execution. The Task Manager cooperates with the Agent Management System where all the objects related with the agent platform are located, including the agent generation mechanism, the communication objects and the proxies to Grasshopper objects. We used Grasshopper [7] for the prototype that we have developed.

The *User Authentication & Authorization System* is responsible for all security and access control issues in our prototype. The authentication process is implemented via a miniLDAP authorization mechanism (downscaled implementation of the Lightweight Directory Access Protocol), which grants or denies permission to access a digital object (data or resource). In order to define these permissions, we have adopted a role-based access control mechanism to handle the users and objects. These roles are shared inside the distributed system by using a hierarchical structure, based on trust domains. Authorization is achieved by validating resource access, signed with certificates that were issued by the ARION system.

AN ENVIRONMENTAL E-SCIENCE SCENARIO

There is a broad range of computational applications in this field. Within the ARION project, we have mainly concentrated on three major use cases. The first one deals with the cooperation of a Technical University (providing ocean wave models) with an oceanographic company (providing ocean data and measurements). The second use case involves the collaboration between two research institutes, using our system as a research environment for experimental inter-organizational scientific (computational) workflows. The third use case manages inner-organizational workflow executions, keeping track of relevant information concerning the workflows (number of executions, input files used, parameters used etc.) and presenting final results to possible customers of this company.

In this section we describe the first, and simplest, use case, the “ocean wave statistics scenario”. The purpose of this scenario is the statistical processing of long-term time series of wave parameters on various geographical locations around Europe and the production of basic statistics, histograms and analytic Probability Density Functions (PDFs) for surface wave data. The statistical processing procedures and tools were developed by one project partner, while the calibrated long-term time series were provided by another partner.

Initially, the user chooses a specific type of Scientific Data, by browsing a hyperbolic tree that contains our ontology. An applet (shown in Fig. 1) provides a view of the entire ontology structure or even different facets of the ontology. By selecting a specific class (concept) in the schema, the user is able to see all instances (i.e. data sets) and attributes of the selected class, as well as metadata information regarding a specific instance of the class. It is also possible to create or edit metadata instances of a selected class.

After that, some predefined queries concerning the type of data chosen are presented to facilitate the user. At this point the search can be further refined by specifying more precise search criteria. Possible queries in this specific use case are:

- Search for wave data in a specific geographical region (e.g. Mediterranean sea, Atlantic ocean, Baltic sea etc)
- Search for specific wave parameters (e.g. significant wave height, mean wave period, mean wave direction etc)
- Search for wave data in a specific time period.
- Combination of the above.

When a query is executed the matching data sets are returned as a set of references. By following such a reference, the user can view all the metadata information related to the selected data set. The query results can also be further refined by selecting one or more specific attributes to be shown. From the metadata, the user is able to access the raw data and/or visualize them. In case no results are returned or the user is not satisfied with the existing data sets, the user can browse for workflows that can produce the type of scientific data required.

The tasks performed during a workflow execution, shown in Fig. 2, are described below:

- Selection of the geographical area of interest.
- Visualization of the available data points and the corresponding geographical area, by means of a geographic GUI.
- Selection of a number of wave parameters and data points for which statistical information are requested.
- Selection of the time period of interest.
- Selection of the appropriate partition(s) for the calculation of histograms.
- Retrieval of the appropriate time series on the corresponding wave parameters.
- Statistical processing and production of the corresponding histograms (univariate, bivariate, etc).
- Fitting of analytic (known) probability models.
- Production and visualization of the results both in ASCII and graphic format.




Fig. 1 Hyperbolic tree view of our ontology

User interaction is required for the selection of the geographic location of interest, the selection of specific points and the specification of statistical parameters concerning the data requested. This interaction is accomplished through a map applet shown in Fig. 3.

Some typical results of such a workflow execution (basic statistics, histograms and analytic PDFs for surface wave data) are shown in Fig. 4. Metadata concerning these results are automatically stored into the system.




Fig. 2: Workflow diagram.




Fig. 3: Map applet GUI for ocean wave statistics scenario.




Fig. 4: Workflow execution and graphical results for the ocean wave statistics scenario.

OPPORTUNITIES FOR SEMANTIC WEB TECHNOLOGIES

In the previous scenario, metadata is needed to describe how to run the scenario, to describe the input to be provided by the user and how to present the output. Metadata information is also used in the various stages of the workflow being executed to ensure that valid input is fed to each stage. Consequently, the metadata technologies offered by the Semantic Web are applicable and relevant in creating a metadata-based middleware. Resources can be discovered, allocated and disbanded dynamically and transparently to the user.

We use Semantic Web technologies to represent the syntactic data types of scientific objects using XML Schema data types, and also represent domain ontologies. Ontologies are widely recognised as a useful mechanism for the classification of metadata for various resources. However, such annotations would be of limited value to automated processes unless they share a common understanding as to their meaning of terms in a given domain. We usually attribute the notion of ontology to the “specification of a conceptualization” [8] - that is, defined terms and relationships between them, usually in some formal and usually machine-readable manner.

An ontology contains a set of classes and each class has an associated set of properties. Each property has a range indicating a restriction on the values the property can take. An ontology relates more specific concepts to more general ones. Such links are used to organize concepts into a hierarchy or some other partial ordering, which is used for storing information at appropriate levels of generality and automatically making it available to more specific concepts by means of a mechanism of inheritance. Ontologies may be grouped into the following three areas, according to their role: to assist in communication between people, to achieve interoperability among computer systems, or to improve the process and/or quality of engineering software systems [17].

The general approach to semantic integration has been to map terms and concepts onto a shared ontology. A shared scientific ontology ensures total integration, though constructing such an ontology may be costly and time-consuming. The creation of agreed-upon metadata and ontologies is usually a community-led process, describing practices established in that community, and also needs a formal and shared representation.

In ARION, we were interested in the semantics of scientific collections. We designed an environmental ontology comprising mainly of two different facets, describing the scientific data and scientific models respectively (see Fig. 5). This way we can provide modularization of a potentially large monolithic ontology. A facet-based engineering of an ontology scales well with large scientific ontologies. The ontology definition contains an “is-a” hierarchy of relevant domain concepts, relationships between concepts and properties of concepts.




Fig. 5: The structure of our environmental ontology.

Considerable efforts were made to develop technologies for asserting facts about resources, and a common language for expressing metadata and knowledge embodied by ontologies; notably RDFS [2], DAML+OIL [10] and OWL [15].

In our implementation, scientific collections are semantically described by concepts (classes) that are defined in an ontology constructed with RDFS, while the representation of resource metadata is realized by RDF. Objects, classes, and properties can be described using a standardized syntax (XML) and a standard set of modeling primitives like instance-of and subclass-of relationships. The expressive capabilities of RDF and RDF Schema suffice for the purposes of ARION and are used as the basis for modeling our domain of knowledge. In particular, metadata description is ontology-driven, in the sense that the construction of the metadata information is carried out in a top-down fashion, by populating a given ontology, rather than in a bottom-up fashion. Every scientific object (data set or model) is described by a collection of attributes (properties), inherited from its parent-class or native to the specific object.

Metadata would also be used within a grid computing infrastructure, for example, at the data/computational layer (classification of computational and data resources, performance metrics, job control, management of resources), at the information layer (schema integration, workflow descriptions), at the knowledge layer (problem solving selections), governance of the grid (access rights, personal profiles, security groupings), accounting infrastructure (billing, computational economy).

Most of the activities in a day-to-day scientific work involve many tasks that need to be executed in a specific order. A workflow trail records how a collection of application-resources is orchestrated so it can be replicated or replayed, or act as evidence, and is represented by a workflow language, annotated with RDF-based metadata.

Scientific knowledge can be replicated and archived for safe-keeping. It is essential to be able to recall a snapshot of the state of understanding at a point in time in order to justify a scientific view held at that time. As data collections and analytical applications evolve, keeping track of the impact of changes is difficult. Storing workflows and the associated metadata is a step forward towards that goal.

The ARION middleware [11] provides: (a) tools to create and publish computational workflows according to an XML based workflow specification, using a customized version of XRL [18], (b) a runtime platform to execute the workflows and (c) the means to store and retrieve workflows.

The manual creation and maintenance of metadata is generally problematic. People are not always in the best position to

create or produce accurate metadata, through circumstance or error. In ARION, workflow applications are designed to include a mechanism to automate the creation and management of metadata.

CONCLUSIONS

The integration of grid technology with ontologies could significantly affect the use of adding value e-services and the ability to extend programs to perform tasks for users more efficiently and with less human intervention. ARION provides a middleware architecture realizing e-science in the domain of environmental systems. It is capable of integrating collections of scientific datasets, including simulation models and associated tools for statistical analysis and dataset visualization. These collections represent application software in several scientific domains, they reside in geographically disperse organizations and constitute the system content. It also actively supports on-demand scientific data processing workflows. Recent standards and their software implementations have made this possible. More advanced e-services, which depend on the scientific content of the system, can be built upon this infrastructure, such as decision making and/or policy support using various information brokering techniques.

It is clear that, as e-science technologies continue to evolve and mature, in the next few years, an increasing number of applications and large-scale systems harnessing the vast potential offered by these technologies will be brought forward.

REFERENCES

- [1] S. Alexaki, V. Christophides, G. Karvounarakis, D. Plexousakis and K. Tolle, "The ICS-FORTH RDFSuite: Managing Voluminous RDF Description Bases", In *Proc. of 2nd International Workshop on the Semantic Web, WWW10*, 2001.
- [2] D. Brickley and R.V. Guha, *Resource description framework (RDF) schema specification*. Technical report, W3C Proposed Recommendation. <http://www.w3.org/TR/PR-rdf-schema>, 1999.
- [3] I. Foster and C. Kesselman (eds), *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, 1998.
- [4] I. Foster, C. Kesselman and S. Tuecke, "The Anatomy of the Grid: Enabling Scalable Virtual Organizations", *International Journal of Supercomputer Applications*, 15(3), 2001.
- [5] I. Foster, C. Kesselman, J. Nick and S. Tuecke, *The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration*, Open Grid Service Infrastructure WG, Global Grid Forum, June 22 2002.
- [6] *Global Grid Forum*, <http://www.gridforum.org>
- [7] *Grasshopper 2 Agent Platform*, <http://www.grasshopper.de>
- [8] T. R. Gruber, "A Translation Approach to Portable Ontology Specifications". In *Knowledge Acquisition*. vol. 6, no. 2, pp. 199-221, 1993.
- [9] J. Hendler, "Agents and the Semantic Web", *IEEE Intelligent Systems Journal*, Vol. 16, No. 2, pp. 30-37, March/April 2001.
- [10] I. Horrocks, "DAML+OIL: a reasonable web ontology language", in *Proc. of EDBT 2002*, March 2002.
- [11] C. Houstis, S. Lalis, M. Pitikakis, G. Vasilakis, K. Kritikos, A. Smardas, "A grid service-based infrastructure for accessing scientific collections: The case of the ARION system", *The International Journal of High Performance Computing Applications*, Vol. 17, No 3, pp.269-280, Fall 2003
- [12] C. Houstis, S. Lalis, V. Christophides, D. Plexousakis, M. Vavalis, M. Pitikakis, K. Kritikos, A. Smardas and Ch. Gikas, "A Service infrastructure for e-Science: the case of the ARION system", in *Proc. of E-Services and the Semantic Web workshop (ESSW2002)*, Toronto, Canada, May 27-28, 2002, Springer, LNCS 2512, pp. 175-187.
- [13] G. Karvounarakis, S. Alexaki, V. Christophides, D. Plexousakis and M. Scholl, "RQL: A Declarative Query Language for RDF", in *Proc. of Eleventh International World Wide Web Conference (WWW2002)*, Honolulu, Hawaii, USA, May 7-11, 2002, pp. 592-603.
- [14] O. Lassila and R. R. Swick, *Resource description framework (RDF) model and syntax specification*. Technical report, W3C Recommendation. <http://www.w3.org/TR/REC-rdf-syntax>, 1999.
- [15] *OWL Web Ontology Language 1.0 Reference*, <http://www.w3.org/2001/sw/WebOnt/>
- [16] *The Semantic Grid*, <http://www.semanticgrid.org>
- [17] M. Uschold, M. Healy, K. Willamson, P. Clark and S. Woods, "Ontology reuse and application". In Guarino, N., (Ed.), *Formal Ontology in Information Systems*, Trento, Italy, pp. 179-192, 1998.
- [18] W.M.P. Van der Aalst and A. Kumar, "XML Based Schema Definition for Support of Inter-Organizational Workflow", *University of Colorado and University of Eindhoven report*, 2001.

ESA RSSD Science Archives User Interfaces and Inter-Operability Systems

5-7 October 2004, ESA/ESRIN, Frascati, Italy

Christophe Arviset⁽¹⁾, John Dowson⁽¹⁾, Jose Hernández⁽¹⁾, Iñaki Ortiz⁽¹⁾,
Pedro Osuna⁽¹⁾, Jesus Salgado⁽¹⁾, Guillermo San Miguel⁽¹⁾, Aurèle Venet⁽¹⁾

⁽¹⁾ESA, European Space Astronomy Centre
Urb. Villafraanca del Castillo
PoBox 02020 Madrid, Spain

Christophe.Arviset@esa.int, John.Dowson@esa.int, Jose.Hernandez@esa.int, Inaki.Ortiz@esa.int,
Pedro.Osuna@esa.int, Jesus.Juan.Salgado@esa.int, Guillermo.San_Miguel@esa.int, Aurele.Venet@esa.int

INTRODUCTION

Within its overall responsibility for scientific operations development and execution for astronomy missions, the Science Operations and Data Systems Division of the European Space Agency's Research and Scientific Support Department (RSSD) carries the role of developing science archives, providing easy access to mission data for the scientific community worldwide.

ESA centre in Spain near Madrid, previously called VILSPA, has now been renamed ESAC for European Space Astronomy Centre. ESAC will host all ESA astronomy missions Science Operations Centres as well as all ESA astronomy and planetary missions' archives. In particular that includes the ISO Data Archive, the XMM-Newton Science Archive, ESA internal INTEGRAL Science Data Archive and the Planetary Science Archive for Rosetta, MARS EXPRESS, SMART-1, GIOTTO, and all ESA Planetary missions.

All ESA science archives developed in ESAC use the same open 3-tier architecture. This architecture, coupled with the use of JAVA and XML allows to have powerful and user friendly user interfaces which offer many functionalities to the scientific communities for querying, seeing and retrieving on-line the data stored in the archives.

Furthermore, some Archive Inter-Operability Systems have also been developed for some of these archives to allow users, remote applications or archives to directly query and retrieve data products from these archives in an automatic manner, without requiring human intervention. Initially developed to fit project specific needs, these inter-operability tools have been adapted to fit the worldwide Virtual Observatory for Astronomy initiative where inter-operability standards have been defined.

ESA SCIENCE ARCHIVES AT ESAC

Archives Projects at ESAC

ESAC is becoming the centre for all ESA Astronomy missions. That will include the Science Operations Centre, some Mission Operations Centre, as well as the scientific archives for data storage and data distribution to the Scientific Community.

The ISO Data Archive (IDA) has been opened to the community since December 1998 and offers a state-of-the-art archive facility to access ISO data products and auxiliary files. It is accessible at <http://www.iso.vilspa.esa.es/ida/>. It contains all the ISO raw and fully processed, science and calibration data as well as all ancillary data (engineering, uplink and downlink data) for a total of about 400 GBytes stored on magnetic disks. Through a powerful and user friendly Java User Interface, over 1200 registered users have already downloaded the archive content more than eight times since it was opened to the public.

The XMM-Newton Science Archive (XSA) has been available to the scientific community since mid April 2002 at <http://xmm.vilspa.esa.es/xsa/>. It contains all the XMM-Newton raw and fully processed, science and calibration data as well as some ancillary data for an expected total of about 2 TBytes stored on magnetic disks. It also contains the XMM-Newton source catalogue, with powerful search criteria. Based on the IDA architecture and code, it offers similar data

query and retrieval (FTP and CDROM) facility through a user friendly User Interface. Proprietary data is accessible to the observation owner only while public data has unrestricted access. Furthermore, the XSA offers an on-the-fly reprocessing facility to allow the users the get the data calibrated with the latest version of the pipeline processing software.

The INTEGRAL SOC Data Archive (ISDA) is for the time being an internal archive for INTEGRAL Science Operations purposes. It has been available since December 2004. It contains all the INTEGRAL raw and fully processed data for a current size of 2 TBytes stores on magnetic disks. Based on the XSA, it offers same type of functionalities, including various level of searches as per INTEGRAL data organization (observations, observations data sets, science window,).

The Planetary Science Archive (PSA) is planned to hold all ESA Planetary missions datasets. The first version was released in March 2004 at <http://www.rssd.esa.int/PSA> with GIOTTO data and some ROSETTA calibration data. In the course of 2004, Mars Express, Rosetta and Smart-1 data will be included, later followed by Huygens. All future ESA planetary missions (Venus Express, BepiColombo) data will also be part of the PSA. Also based on the XSA architecture, the PSA is the central place for all ESA planetary missions and offers similar power search, view and retrieving capabilities. Furthermore, a PSA Validation Tool has been developed to allow the Mission Instruments Teams to validate their Datasets against the PDS standards before they are ingested into the PSA.

One Team, Several Projects

Since 2002, a dedicated Archive Development Team has been created at ESAC. It provides support to various archive projects. Each member of the group is not dedicated to one project in particular, but is working on a specific sub-system across all the archive projects. There are some Database engineers, some experts on User Interface and Middle-Tier and others on Data Storage, Data Distribution and Interoperability aspects.

This ensures permanent knowledge transfer from one project to another as people are working on several projects at the same time. Faster development is then made possible through this efficient manpower utilization. This approach also permits a more flexible manpower allocation. Some “old” projects do not have enough money to maintain several software experts on distinct fields. But by sharing them with new projects, “old” project can afford to pay for fraction of expertise, while new projects benefit from the previous expertise and get development made faster and cheaper. Furthermore, people remain strongly motivated by participating to new projects and discovering new technology, while at the same time still providing support part-time to some less challenging past projects.

Nonetheless, scheduling people’s activities accross all these projects may result difficult as each individual project has its own milestones which may conflict between each others. Furthermore, projects are paying according to the amount of mandays they are using. Regular and clear monitoring of the work per person and per project is of vital importance to make sure each project really gets the amount of effort it is paying for. Team members report on a weekly basis their tasks and associated effort per project. Monthly, the team leader provide report per project, indicating the list of tasks completed and the corresponding effort spent. This transparency is a key factor to confidence from project towards the Archive Development Team.

Stay Close to the User

Each archive project has nominated an Archive Scientist who is the main interface between the scientific community and the Archive Development Team performing the development. The Archive Scientist is an expert in that science field and has the role of consolidating the users requirements. Regular interactions with the developers ensure mutual comprehension and good cooperation to make sure the final archive meet users’ expectations.

An Open -Tier Architecture

All the archives are based on an open 3-tier architecture described in Figure 1.




Fig. 1. ESAC Archives Open 3-tier Architecture

The main goal of this architecture is to separate the data from the presentation, which allows a more modular and flexible development. Data are saved on UNIX magnetic disks for direct access. From the data products, metadata is extracted and put in a Relational Data Base, S BASE. The middle tier, also called the Business Logic, provide a transparent access to the data products and to the metadata. This key layer has been developed in Java and XML and resides on the archive server. Normal use of the archive is for the end users to access them through the standard Java applet user interface. Furthermore, remote applications and other archives can also direct access the data and the metadata, bypassing the standard User Interface, by speaking to the Business Logic. This has also been adapted to the standards adopted by the Virtual Observatory (VO) for the XSA and the IDA to be fully part of these interoperability initiatives.

Re-using same architecture from previous projects brings the advantage that less effort is required in designing new projects. Furthermore, one is ensured that the architecture is robust as it has been implemented and used successfully already. Re-using same code obviously bring the overall development costs down in the short term. In the long term, it also ensures cheaper maintenance as fixes or improvement are applied to all archives at the same time. The variety of end users around the different scientific communities makes sure that problems are found (and later on fixed!) and makes the archive more robust. This approach should allow the much-quoted paradigm: “**Better, Faster, Cheaper**” to be actually achieved in archive development and maintenance

ARCHIVES USER INTERFACES

Querying the Archives

All Archives can be access by a Java Applet through any browser. They offer powerful query facilities (see Figure 2) with hundreds of search criteria organized by panels organized by area, which can be open or closed by the user to help clarity. By entering values for the search parameters, the corresponding SQL statement is built by the interface, and the user can also view and edit such statement if he is expert enough. Value lists are updated according to other relevant criteria already selected. Sort facilities for the result are also possible for the user. Contextual on-line help is provided for each criteria on the interface.




Fig. 2. ESAC Archives User Interfaces: Query Panels

Viewing the Results

Once the user has selected his criteria, he can execute his query to obtain a result (see Figure 3). The result displays offers the list of items matching the search criteria in a hierarchal form depending of each project (eg for XSA, observations, exposures and sources) and reflected by the color scheme. The number of items displayed as well as the parameters for each of them is configurable by the user and navigation buttons are provided for long lists. By default, all data items are displayed but only public data have a preview images and can be retrieved instantly by the one-click direct download button. Similarly, contextual on-line help is provided for each parameter displayed on the interface.



Fig. 3. ESAC Archives User Interfaces: Results Panels

Retrieving the Data

Once logged in, user can retrieve data in a fully automatic way, without any human intervention. As explained before, this can be performed through a one-click direct download for a specific items listed in the catalogue. Furthermore, this can also be done through a shopping basket mechanism. When viewing results, the user can move items of the catalogue to the shopping basket, perform more queries, move more items, until he is happy with his selection. The full data package will be prepared for download on the FTP area and the user informed of its availability by email. That process usually takes a few minutes as all the data is stored on-line on magnetic disks. Various pre-defined levels of processing can be specified for the retrieval, even some custom datasets, where the user can decide the individual file types he is interested in.

Use of Java and XML

In general, Java provided the multi-platform support, having the same code for all operating systems, realizing the dream of many developers: *write once, run anywhere*. It provides easy deployment on user’s browser while producing a clean and compact user interface that looks like an application, as opposed to the other web approach based on cgi-bin scripts.

One can now go one step further by using XML to achieve multi-project support. In conjunction with the multi-platform capability of Java, the use of configurable XML files allows the same code to be used for various archives to reach *write once, use anywhere* capability.

Within all the archives, XML is used in various ways. First, semi-dynamic XML configurable files are used to describe the appearance and associated action for each element (button, input field, etc..) of the user interface query panels. Similarly, other XML files define all the graphic elements to be displayed as a result of a query. Second, the database is described with XML files which allows the Business Logic to make transparent access to the metadata, regardless of the RDBMS used (S BASE, ORACLE, ...). Third, data distribution request is made via a XML file, which is then mapped to the data products files located on the data storage.




Fig. 4: Re-usability of JAVA code to build the Archive User Interfaces

JAVA and XML have allowed design and code re-use between archive projects (see Figure 4), with minimum change, just modifying the XML configuration file. This resulted in having 95% of the Java code identical for the User Interface and Business Logic system between archives projects.

Same Look and Feel of User Interfaces

Re-use of the same code for the IDA and the XSA user interfaces also offers similar Look & Feel. The user then immediately recognizes that he/she is accessing an ESA archive, leading to a clear ESA image within the community.

INTER-OPERABILITY : STEPS PROCESS

Making ESA science archives inter-operable with other external archives and applications has been a 3 steps process as we have gained experience to make the inter-operability system more and more flexible.

Step 1: Specific Inter-Operability Services

At the beginning, external archives or applications needed to know our archive content (eg: list of observations) to be able to access our two services:

- Postcard Server: service accessible via HTML which returns the preview image of the observations, embedded in an HTML page with links to relevant documentation. Example of such service can be found at:
<http://xsa.vilspa.esa.es:8080/aio/jsp/createPostcards.jsp?obsno=0112570601>
- Product Server: service accessible via HTML which returns via FTP the data associated with the observation. Example of such service can be found at:
<http://pma.iso.vilspa.esa.es:8080/aio/jsp/product.jsp?obsno=57600101>

These services have already been in use for years by external archives like CDS Vizier and Aladin (France), ADS, IRSA and HEASARC (USA), but they present the inconvenience requiring regular update of the archive metadata content on the remote site and offer limited scriptable capabilities.

Step 2: Archive Inter-Operability System (AIO)

A more generic Archive Inter-Operability (AIO) system was developed to circumvent the previous described drawbacks. The AIO system has been built as a complete new layer on top of the existing archive business logic to offer more services as required by the users:

- Existing Postcard and Product Server in HTML modes as described in previous section
- Access to archive metadata content via an XML scriptable interface
- Access to product server in a socket mode (XML or command line request) which returns data through the socket

This new AIO system provides powerful scriptable interfaces for internal and external users and has been extensively used by Scientists for automatic calibration and monitoring purposes.

Full documentation for the ISO and XMM-Newton AIO systems can be found at:

- <http://pma.iso.vilspa.esa.es:8080/aio/doc/>
- <http://xsa.vilspa.esa.es:8080/aio/doc/>

Step 3: VO Compliant Archive Inter-Operability System

The Virtual Observatory is a world-wide initiative in astronomy which aims is to allow astronomers to perform new science by providing them with a “federation of astronomical archives and databases around the world, together with analysis tools and computational services, all linked into an integrated facility”. Some Inter-Operability standards have then been defined, eg VOTable (XML based) for metadata transport and SIAP (Simple Image Access Protocol) and SSAP (Simple Spectra Access Protocol) for data archive query (metadata) and data retrieval (products). These protocols are still under definition and evolution is to be expected. Therefore, there was a need to keep our archives inter-operability system flexible so we could adapt easily with existing and future VO standards (see Figure 5).




Figure 5: ESAC Archives and Virtual Observatory

By building “translation layers” using XML transformation rules (see Figure 6), we could easily adapt our existing AIO systems to these new VO standards. VO standards metadata query and data requests are translated into our AIO metadata and data request to be passed to our business logic. Similarly, metadata results and data products were translated from our system towards the VO standards.




Figure 6: AIO and Translation Layers

CONCLUSION

ESA European Space Astronomy Centre (ESAC), located near Madrid in Spain will host all ESA astronomy and planetary missions’ archives. All these archives are developed by the same team which ensures knowledge transfer between projects and allow re-usability and therefore significant savings in development, operations and maintenance costs.

Based on the open 3-tier architecture and developed in JAVA and XML, the archives’ user interfaces offer user friendly and powerful features for querying, viewing and retrieving data, all in a fully automated way.

This modular architecture has also allowed the archives to inter-operate easily with external archives and applications, following inter-operability standards laid out by the astronomical Virtual Observatory in initiatives worldwide.

REFERENCES

- [1] C.Arviset, “ESA RSSD Scientific Archives”, 8th ICATPP Conference on Astroparticle, Particle, Space Physics, Detectors and Medical Physics Applications, November 2003, *World Scientific*, in press.
- [2] P.Osuna, J.Hernández, C.Arviset, J.Dowson, “XML in the XMM-Newton Science Archive”, XSPACE 2001 Conference
- [3] C.Arviset, P.Osuna, J.Salgado, “ESA ISO and XMM-Newton Archives Inter-Operability and VO Services”, Astronomical Data Analysis Software & Systems Conference, October 2003, *ASP Conference Series*, in press

Discover, browse and access ocean data on internet

Sophie Baudel⁽¹⁾, Frédérique Blanc⁽¹⁾, Tony Jolibois⁽²⁾, Vinca Rosmorduc⁽¹⁾

⁽¹⁾CLIS, 8-10 rue Hermes, 31526 Ramonville-St-Agne Cedex, France
Email : baudel@cls.fr, blanc@cls.fr, rosmorduc@cls.fr

⁽²⁾COFRAMI, Le pré Catelan, 97 rue Riquet, 31000 Toulouse, France
Email : jolibois@cls.fr

INTRODUCTION

The Products and Services (P&S) department in the Space Oceanography Division at CLS is in charge of diffusing and promoting altimetry and operational oceanography data. P&S is so involved in Aviso (satellite altimetry), Mercator ocean forecasting system, and European Mersea ocean portal.

The activity began with altimetry data [1] by diffusing data on hard supports which still remains a useful mean to many users who are not supplied with high bandwidths.

In the 95th, first internet altimetry data transfer by FTP began with very slow data flows due to small bandwidths. This diffusion mode took a much better place in the 1999-2000 with the coming of high speed transfers.

Thanks to the coming of high bandwidths and simultaneously client/server architectures, in the 1999-2000, the next step for us consisted in the implementation of the Opendap server opensource software and its Life Access Server (LAS) user interface.

Our purpose here is to describe the use of this technology in our P&S department for two types of data : operational oceanography data (Mercator French project and European Mersea project) and altimetry (Aviso portal).

OPENDAP/LAS TECHNOLOGY OVERVIEW

Opendap

Opendap (Open source Project for a Network Data Access Protocol), formerly known as DODS is at the same time :

a requests and answers transfer web protocol based on HTTP,
a set of server applications (java servlets or CGI applications) allowing the user to extract georeferenced data by varied client softwares as Ferret, Matlab, IDL [2].

Opendap allows the user to extract the data he is interested in and only this data, avoiding him to download full information files. Opendap allows particularly to extract :

- a geographic area,
- a period time,
- an oceanic variable,
- an output format.

Several Opendap servers exist [3]. The one we have implemented in CLS is the DODS Catalog/Aggregation Server (AS) relevant for our time-evolving data type. It is a Java servlet. It's main purpose is to allow the "aggregation" of multiple datasets into one DODS dataset. It creates virtual datasets from NetCDF files. It uses THREDDS catalogs to specify what datasets it serves, and how to aggregate them.

The Opendap can be used directly via one’s own software client. The user can browse the Opendap catalogue and get informations on the data sets. Fig. 1. and Fig. 2. show the two AS catalogues such as they appear on an Internet browser, that we have implemented for Aviso altimetry and Mercator operational oceanography data respectively.




Fig. 1. Aviso Opendap data catalogue (<http://opendap.aviso.oceanobs.com>, free access)




Fig. 2. Mercator Opendap data catalogue (<http://opendap.mercator-ocean.fr>, authorized access)

LAS

LAS is an Opendap data access highly configurable web server. Its default visualisation application is Ferret, an opensource software dedicated to gridded geo-referenced scientific data. Ferret is an Opendap client.

Whereas Opendap direct use requires from the user to have the knowledge and possession of a client software such as IDL or Matlab, LAS allows a not very well-informed user to visualise the data very easily. Some examples of what he can do with the LAS are given below :

- map data with on-the-fly graphics,
- request subsets of variables in a choice of different file formats,
- compare (difference) variables from distributed locations,
- compute basic statistics (average or variance).

The special feature of the LAS consists in the facility for unify in a single vision the access to multiple types of data from distributed data sources. The LAS can make requests to different remote Opendap servers. This enables to make comparisons or statistics upon several different data types.

APPLICATION TO OCEAN MODELS DATA (MERSEA/MERCATOR)

Mercator-Ocean is a French operational oceanography Center which distributes its products by three means :

- Weekly web bulletins : a set of GIF images or animations which give an overview of the main oceanic analysed and forecast variables of the several Mercator prototypes. These images are weekly updated (historical datasets remain on line) and available on the Mercator web site [4].
- CNES/MAPS server which enables remote access to numerical NetCDF or original raw output files. The identified and authorised user select the data he wants. He can retrieve them by direct downloading or on hard supports.
- LAS/Opendap servers as part of Mercator Mersea-strandI [5] contribution.

The Mersea LAS [6] unifies four other ocean models (Fig. 3) : the UK Met Office FOAM Atlantic and Mediterranean model, the Norwegian NERSC Atlantic Topaz model, the Italian INGV Mediterranean model and the US Hycom model.




Fig. 3 Mersea strandI LAS interface : presentation of the four Atlantic models datasets

The main difficulty of this LAS implementation has lied in the ocean model metrics definition and a common file format adoption which forced the four model teams to produce the same datasets in the same formats. Notice that this was a pioneer approach and that it has been adopted by Godae (Global Ocean Data Assimilation Experiment) standards.

We show now some useful applications of the LAS server, such as intercomparison facilities (Fig. 4), a temperature section (35 N) across the Atlantic (Fig. 5), an hovmoller (xt) diagram showing variations of salinity between 1°W and 5°E at 37.5°N (Gibraltar strait area) issued from the italian MFS operational model, a surface current one-week forecast in the Galicia area (Fig. 6) as seen by the Mercator high resolution model for the 16/07/04.




Fig. 4. An example of a difference field between two models outputs : Mercator and Foam (sea surface height, North of Spain area, 23/06/04)




Fig. 5. An example of temperature section from surface to 1500m at 35°N across the Atlantic as computed by the Mercator French model for the 01/07/04




Fig. 6. An example of an hovmoller (xt) diagram showing the time evolution of 5m-salinity between 1°W and 5°E as computed by MFS Italian Mediterranean model




Fig. 7. An example of a surface current forecast representation in the area of Galicia, as seen by the Mercator high resolution prototype

APPLICATION TO ALTIMETRY DATA (AVISO)

Aviso is the CLS service which distributes altimetry products since 1993 [1]. The Aviso Opendap/LAS server has been implemented in 2003 simultaneously with the Mersea/Mercator Opendap/ LAS one. The Aviso LAS distributes several Ssalto/Duacs altimetry products such as delayed and near-real time mean sea level anomaly, absolute dynamic topography, absolute geostrophic velocities, gridded significant wave height and gridded wind speed modulus. These data are issued from a set of 5 satellites among them four are active : GFO, Jason, Envisat, Topex/Poseidon. Merged products issued from the different satellites data combination are also available.

CONCLUSION AND PERSPECTIVES

Opendap technologie allows to create several servers to distribute the same data to different user profiles without duplicating the data. In this way, CLS is implementing several Mercator Opendap servers which correspond each to a specific use. Several user profiles have been defined such as research profile (all the variables, all the levels, full resolution), large public profile (a few basic oceanic data at a reduced resolution), race-sailing profile (only surface currents data). Other perspectives deal with the implementation of a Map Server, a GIS opensource server which will communicate with the Opendap server. The Map server will be able to manipulate simultaneously raster and vector remote data. The aim is to construct a full complete web oceanic data distribution services. The projects in which we are involved allow us to progress towards that.

REFERENCES

- [1] V. Rosmorduc, "12 years of user services for ocean topography users : Aviso experience and lesson learned", this issue.
- [2] Available Opendap client software : <http://www.opendap.org/faq/whatClients.html>
- [3] Available Opendap servers : <http://www.opendap.org/faq/whatServers.html>
- [4] Mercator weekly bulletins website : <http://bulletin.mercator-ocean.fr>
- [5] Mersea strand1 products website : <http://www.mersea.eu.org>
- [6] Mersea strand1 LAS : <http://las.mersea.eu.org/las/servlets/dataset> (authorized access)

On Using Grid Computing Technology with the OAIS Reference Model to Provide Persistent Archive Environments and Value -Enhancement for NASA Earth Science Data

PV-2004, "Ensuring the Long -Term Preservation and Adding Value to the Scientific and Technical Data"

5-7 October 2004, ESA/ESRIN, Frascati, Italy

Bruce R. Barkstrom, David E. Corder

Affiliation

*Atmospheric Sciences Data Center
MS 157D*

*NASA Langley Research Center
Hampton, VA 23681-2199 USA*

Email: Bruce.R.Barkstrom@nasa.gov, David.E.Corder@nasa.gov

INTRODUCTION

Grid computing technology appears to offer significant advantages in building persistent, federated archives and adding value to existing collections of data. In the environment at the NASA Langley Atmospheric Sciences Data Center (ASDC), as well as many other centers for Earth science data, users suffer from several kinds of difficulties:

1. Data from the current archives are primarily available as large, binary files, organized in a form that may require significant effort on the part of users to master.
2. Significant effort and scientific expertise is required to combine data sets from different instruments to produce interdisciplinary insights into interesting scientific phenomena.
3. Conventional approaches to data throughput may limit the ability of scientists to interact with very large volume data sets of the type needed for climate research.

Grid computing technology can help this situation by increasing data throughput in and between archives. In particular, it can help users in several distinct ways:

1. Creating on-the-fly, non-image subsets and visualizations for large collections of files from the same data set version.
2. Efficiently and automatically reformatting data files into formats more useful to the user community.
3. Combining data from coincident and synchronous observations for large collections of files - including the ability to combine data held at several different data centers (e.g., TMI and CERES - TMI data at Ames Research Center and CERES data at Langley Research Center). As a specific example, we have done some small scale exploration of how to locate mesoscale convective systems in TMI data, find the equivalent footprints of CERES, and then create a federated catalog that could extract instances of these systems efficiently.
4. Putting some of the Level 3 data into databases to produce regional time series that the archive can supply to users through a "web-services" (or, perhaps more appropriately, "grid services") approach.

In addition, Grid computing technology appears likely to reduce operating costs for persistent archives by providing tools for automating substantial portions of an archive's activities.

We discuss some of our experiences in grid computing with commodity-based cluster computers, as well as progress in working on standards for implementing the Open Archive Information Systems (OAIS) Reference Model (RM) within the Persistent Environments Research Group of the Global Grid Forum.

CURRENT CHALLENGES TO USERS OF NASA EARTH SCIENCE DATA

In an ideal world, using Earth science data would be as simple as obtaining a suitably organized ASCII file from a data center, reading it with a spreadsheet program, manipulating it within that tool, and then using the spreadsheet analysis to reach important conclusions for scientific work or to reach important decisions. In the case of the Solar Surface Energy (SSE) data from ASDC, that vision is practiced now (see http://eosweb.larc.nasa.gov/PRODOCS/sse/table_sse.html).

Effective Cloud Height (Terra-FM2 SSF)




Fig. 1. Effective Cloud Height from the CERES Flight Model 2 Instrument on the Terra Spacecraft. These data are from the SSF data product from Hour 2 (UT) on November 15, 2000, with the instrument operating in Rotating Azimuth Scan Mode.

Expert Knowledge Requirements

However, the SSE data are an unusual case – these data are already in the form of statistical summaries of geographically-gridded averages of a climatology. In addition, the uses of these data are clearly enough defined that it has been easy to work with Canadian agencies to create a spreadsheet template which can ingest an ASCII data file without further work.

The bulk of the data in ASDC are in the form of large binary files that contain moderately complex data structures. For example, Level 2 data from the investigation of Clouds and the Earth's Radiant Energy System (CERES) are stored in 200 MB files using Hierarchical Data Format Version 4 (HDF4). Similarly, data from the Multi-Angle Imaging SpectroRadiometer (MISR) have file sizes of 500 MB and also use HDF4. Some of the data in ASDC can be interpreted as images, which allows a user to ingest a file into commercial software for display. However, much of the data are not images. Solar constant data from the Active Cavity Radiometer, record time series of solar irradiance. Data from the Stratospheric Aerosol and Gas Experiment (SAGE) are usually monthly averages of latitudinal variations of vertical profiles of aerosol and gas concentrations. An accurate mental image of the MISR data should probably incorporate the notion that they are three-dimensional movies in nine colors. Level 2 data from CERES are most accurately portrayed as collections of tilted vertical columns of radiation fluxes and cloud properties – with about half the data taken in a mode that has no image topology. Fig. 1, a rendering of a small portion of a swath of CERES data taken from an instrument in Rotating Azimuth Plane scan mode, illustrates this complexity.

It often seems that Earth science data are synonymous with imagery, where features of interest are immediately visible to the untrained eye. However, Fig. 1 should serve as a reminder that scientific data can be much less obvious. Scientists designing instruments and algorithms go to great lengths to extract quantitative information from the data obtained from space-borne instruments. To properly interpret much of the Earth science data in NASA and other archives, users require many skills: knowledge of orbital sampling geometry, the physics of instrument calibrations [1], use of radiative transfer and inversion theory, incorporation of temporal interpolation that allows for complex patterns of variability, and the statistics of uncertainty – to name a few. The teams that have created the data products in the archives have invested thousands of person-years of time into developing and validating the algorithms that produce these data. One measure of this investment is the size of the code used in production. At ASDC, MISR and CERES data are produced with systems whose source code is about one million lines of code each, subdivided into seven to thirteen subsystems. While database experts often describe data product creation as involving “eager” or “lazy” evaluation of data, in practice, Earth science data often requires “expert” evaluation.

An example of the difficulties involved can be illustrated by the problem of spatially interpolating between measurements. Even if the data are in the form of single-band images, the actual measurement values are produced by spatially integrating the underlying radiance field over the instrument field-of-view, weighting with that spectral band's Point Spread Function (PSF):

$$m_i = \text{Integral}_{\text{over FOV}} [dO \text{ PSF}_i(O) I(O)] \quad (1)$$

where m_i is the i 'th measurement in an image, $\text{PSF}_i(O)$ is the Point Spread Function for that measurement, and $I(O)$ is the underlying radiance field [2, 3]. O represents the (complex) geometry that relates satellite position in orbit to the geolocation and viewing geometry, as well as the instrument viewing geometry. A proper interpolation of the radiance field requires creating interpolation functions at high resolution with respect to the sampling mesh and solving the equation

$$K_{ij} I_j = m_i \quad (2)$$

where K_{ij} is the matrix produced by integrating the high-resolution representation of the radiance field and the PSF over the interpolation basis functions [4]. Solving this equation correctly to produce the radiance fields at the high-resolution points, I_j , requires the use of Singular Value Decomposition [5, 6]. For practical work, the solution effectively multiplies noise in the measurements by large values and has the potential of destroying quantitative conclusions for resolution “enhancements” that attempt to provide interpolated values at spatial resolutions smaller than one-half of the sampling distance in the original data. To make this specific, while Geographic Information Systems provide the ability to apply standard interpolation functions (such as bi-cubic splines) to imagery, such interpolation functions may not be appropriate and conclusions drawn from them may or may not be valid. The evaluation of their correctness requires the numerically intensive work of solving equation (2) correctly.

As we move into the future, scientists will clearly require more ability to intercompare and combine data from many different instruments. However, we should not underestimate the difficulty of such an undertaking. We need to remember that it can require several years of contact with a science team and the data it produces to gain sufficient knowledge to avoid serious blunders in interpreting data from a single instrument. Combining data from two instruments requires learning the nuances of the contents of both kinds of data – and then new skills, such as those associated with the estimation of the uncertainty that applies to fused data products.

Data Management Issues

It is one thing to access a file with a single image; it is a very different thing to work interactively with hundreds of thousands of files. One example of this difference arises in thinking of the work involved in doing data mining studies for climate with Level 2 data. Specifically, suppose an individual wanted to study the influence of large-scale cloud systems (such as fronts, clouds associated with low pressure systems, or hurricanes) on the Earth's radiation budget. Because such cloud systems move rapidly, time averaged data are probably not appropriate – important details will be wiped out by the averaging process. It will also be difficult to intercompare the observations from CERES with those from microwave data, say, that are likely to contain useful information about the liquid (or ice) water content of the clouds – although it is important to recognize that cloud water or ice details change on time scales of minutes, so that intercomparison of daily or spatial averages is likely to be extremely noisy. Thus, there is strong justification for intercomparisons of instantaneous data at the original space and time resolution of the data. This creates a requirement for large-scale work with Level 2 data.

For this kind of work, the CERES data product identified as SSF is likely to be a primary source. At the same time, cloud systems have not been identified as such in the CERES processing (as with other NASA Earth observation data). Thus, an individual wanting to conduct this kind of investigation might need to work through twelve instrument years of CERES data. The data volume would be about thirty TB of data in about 105,000 files. It is useful to note that the current archive maintains most of its data in robotic tape silos. For some of the data in the archive, the files on tape are distributed at random through the silo, so that the robotic arms and tape read heads spend $\frac{3}{4}$ of their time in searching and seeking data. This latency reduces the effective throughput at which data can be extracted to about 10 MB/sec. At this rate, it requires the archive about thirty five days to extract the necessary data for this kind of investigation.

It is also important to recognize that this kind of data mining work is intrinsically interactive. The algorithms required for either object identification or object intercomparisons are not simple, nor are they likely to work successfully on their first trial. We must expect that good interpretations will only emerge after a fairly large number of algorithm adjustments, with the adjustments being based on intercomparisons with other data sources and with the statistical properties of the results over global, long-term sets. While much useful work can be done on a small number of files, in the end, scientists are likely to need several iterations through the complete data set in order to produce validated data.

The requirements for interactive work with large volumes of data, sometimes with data at different physical locations, create some interesting logistics problems. The simplistic view is that we can “bring the data to the computation” or “bring the computation to the data”. In practice, the choices may be more complex. An economical solution is needed of optimizing the cost of computation and the cost of data flow (networks) – given that there may be infrastructure constraints on the resources available for both of these processes.

An example of the tradeoffs that arise in the larger-scale domain involves the previous example, but adding in the fact that the network bandwidth currently available from ASDC is about 6.5 MB/sec. At this rate, it would take the data center about 55 days to transmit the total data to the user – assuming that all of that bandwidth could be devoted to serving this single purpose. Increasing this bandwidth involves the politics of multiple organizations and their resources. We do not deal with it here. An alternative is to put the data on tape and ship it to a user. In this case, the cost needs to include the media costs, where the typical tape used for such transfers in the current system holds 32 GB and costs about \$80 per tape. With this set of constraints, the user would need to supply ASDC with about 1,000 tapes to hold this amount of data. The cost of these tapes is currently about \$80,000. A third alternative is to bring the programs into the data center and perform the work there. That approach at least avoids the network and data transfer issues, although it introduces additional work at the data center – and may make it harder for the investigator to complete the work because of the necessity for him or her to interface with the data center.

In practice, the difficulties of moving the data over the net or transferring it on media have forced ASDC to concentrate on the approach of bringing the computation to the data. This approach does have the advantage that the data center can optimize its local environment to deal with the data management issues and provide a reasonable solution to the problem of serving multiple user communities within the resources available. It also avoids the difficulties associated with many members of the user community who have little experience with large-scale data management. On the other hand, it does create a need for appropriate visualization software that will allow the user to see the results of the algorithms and to correct problems. Given the bandwidth limitations already mentioned, care is required in tailoring the visualization approach to the actual interactions between the data center and the individual investigator.

COMMODITY COMPUTING AND GRID TECHNOLOGY

The first element needed to solve these problems is the commoditization of computing hardware, and particularly of Linux clusters based on commodity CPU's and disk storage. Whereas a few years ago, high performance CPU's cost about \$20,000 apiece, the commodity price for equivalent compute power has dropped by a factor of twenty – as long as the job to be done involves computing on a “one-file-per-CPU” basis. Storage for this kind of problem has also dropped in price by a similar factor. This fact has improved our ability to increase throughput by a factor of about 25, dropping the time required to examine and find objects in all thirty TB to about a day and one-half, once the data have been extracted from the grips of tape storage. This element also fits well with recent recommendations of the National Research Council [7].

The second element is Grid computing [8, 9], which increases the freedom of a data center to apply a less centralized architecture to the computing and storage within the center. Grid computing is usually identified with three software components:

1. The Globus Toolkit, which is intended to provide communications capability, particularly including reliable ftp

2. The Storage Resource Broker, which provides the important capability of separating the logical name space for storage from the detailed properties of physical location of the data stored in the archive
3. The Condor scheduling software, which allows flexible scheduling in heterogeneous computing environments

Together, these elements form a powerful synthesis for extending the range of services archives can provide to their user communities. They allow an archive to assume a decentralized, peer-to-peer approach in its internal architecture that increases the flexibility with which the archive can apply its resources to solve problems that are hard to deal with when there are very rigid limitations on how compute and storage elements must interact.

First, these elements make it possible to create on-the-fly, non-image subsets and visualizations for large collections of files from the same data set version. With a commodity CPU Linux architecture, each CPU has enough storage to hold both programs and files. This means that an archive can rapidly change its configuration by deleting old programs and data from an appropriate number of CPUs and loading in new programs and data. By making the programs act like filters that accept one file as input and then create a different kind of file on output, we allow a substantial portion of the compute capability of the archive to be applied to one kind of job and then rapidly redeployed. In addition, the abstraction of the system storage allows the archive to disperse the I/O more appropriately across the system's resources and substantially increase throughput without a large investment in human scheduling. As far as the programs that perform subsetting, object identification, or creation of images from a particular data file, the interface must specify the name of the input file, the name of the output file, and the exception handling behavior to be expected from the program. Thereafter, the executable can be assigned to any CPU that is available, making it much easier to change the profile of the computation without a large overhead for machine reconfiguration.

Second, these elements allow the archive to efficiently and automatically reformat data files into formats more useful to the user community. To a substantial degree, this feature of the commodity cluster architecture is derivable from the first. The output of a reformatting program is another file – and therefore no different as far as the functions of the archive are concerned as the output files from a subsetting job. Only the content of the program has changed. We can even view subsetting and reformatting abstractly as very similar operations that accept the stream of data from the input file and reduce/reorder the data elements to produce the output file.

To put this abstract view another way, the commodity Linux platforms that are Grid enabled can be viewed as abstract, flexible data transformation machines. The archive can create a library of transformation elements (i.e., the programs that execute the transformations). Data transformations are performed by supplying a job scheduling engine, such as Condor, with a list of compute and storage resources that can accept abstractly defined packages of computation. This abstraction allows the scheduler to view the job of content-based data searching or of reformatting as essentially identical – the job of the scheduler is to match the abstract packages of computation with abstract packages of input data and then directing the output stream to storage in the appropriate location. From this point of view, a Linux cluster becomes a flexible compute resource that can be rapidly reconfigured from one kind of job to another with low overhead.

Third, the elements we identified above make it possible to combine data from coincident and synchronous observations for large collections of files - including the ability to combine data held at several different data centers. This capability does not necessarily require archives to transfer large amounts of data. Often, it requires the cooperating archives to provide pointer equivalents that allow a feature identified in one set of data to be identified in another.

As a specific example, in conjunction with colleagues at Ames Research Center, we have done some small scale exploration of how to locate mesoscale convective systems in the TRMM TMI data stored at Ames, find the equivalent footprints of CERES stored at Langley, and then create a federated catalog that could extract instances of these systems efficiently [10]. The Ames researchers developed a data mining program that would identify the systems in TMI data, sending the data and the miner to Glenn Research Center. After the miner found these systems, it provided a convex hull that described the geographic area wherein the system was located. There were thirty seven systems in the one day of TMI data we used, with a typical system area of about 2500 km². The convex hulls are small data objects, so that transmission of a file with these features took a very small amount of time to go from Langley to Langley – particularly compared with the time that would have been required to transmit one data set or the other to a common location. Clearly, creating a well-engineered approach to synchronizing object identification can involve some careful algorithmic work.

In the portion of this example done at Langley, we avoided a two-dimensional search for objects identified in longitude and latitude by recognizing that both the TMI data and the CERES data came from instruments on the same satellite. This meant that with the highly inclined TRMM orbit, we could perform a one-dimensional search in longitude to find CERES footprints whose longitudes fell within the range identified by taking the maximum and minimum longitude of a mesoscale convective system. If the footprint longitude was not in this range, there was no point in searching further to check whether the footprint would fall within the system. The conversion of what might seem to be a two-dimensional search into a one-dimensional one saved us a very large amount of compute time by a factor of 10^{-3} to 10^{-6} .

This experience suggests that we need to develop better approaches to distributed data searches than the conventional one of brute force transmission of data from one location to another. Rather, we need to develop the equivalent to "smart pointers" to data features that will allow an object identified in one data set to identify the same feature in another data set, independent of whether the second data set is located in the same facility or on another continent. By creating catalogs of such pointers, we can markedly expand the value we provide to archive users. Indeed, if we take a very long term view, the value provided by the interpretation of the data in federated archives may exceed the original value of the data, even when measured in the workforce that was originally required to create and validate the data products.

Fourth, the elements we have identified coming from Grid computing can also assist in applying database technology to serving user needs. For example, while we currently supply the gridded and time averaged data in "Level 3 data product" files, we may increase the archive's value to the user community by providing regional time series extracted from databases with pre-specified queries, particularly if the archive uses a "web-services" (or, perhaps more appropriately, "grid services") approach.

The view we are taking might be summarized as moving from an older view of an archive to a newer one. In the old view of an archive, the institution supplied files to users who searched through a metadata catalog for the ones that would fill their needs. In the new view, an active archive provides made-to-order data with annotations and interpretations that markedly increases the value to the user. Commodity Linux clusters and Grid computing software substantially decrease the difficulty of achieving this new view.

PRESERVATION

In the old paradigm of data storage, it was hard to achieve independence from physical devices. In the new one, we use a more abstract view of the objects in the archive and their physical location. In other words, archive middleware, such as the Storage Resource Broker, manages both the logical and the physical location of the data. To a substantial degree (and to the extent that archives can obtain sufficient bandwidth), we are moving toward a situation where we will be less concerned about the physical location of the data, as long as appropriate access rights are guaranteed.

This independence is important because it has the potential to increase the probability that the data in the archive can survive over really long periods of time. As background, it is useful to consider how careful we need to be if we wish to preserve data in the face of technological obsolescence and malicious users of that technology.

Consider, for example, the fact that under current conditions, the probability of serious intrusion of sites connected to the Internet is about 10% per year. If we accept the figure that archives need to preserve data for 200 years, then the probability of survival is about 10^{11} if all the copies of the data are kept on-line and connected. If we can store at least one copy off-line (and probably on media of some sort), then we can expect data migration about once every five years – noting that the primary reason for migration is the model obsolescence of the read and write hardware (not the aging of the storage media). In order to have a probability of data survival greater than 99% over 200 years, the data migration process must have a probability of loss below about 10^{-4} for each migration.

It may be difficult to achieve such low error rates, particularly if we consider that both the hardware and the migration software need to be included in the migration process. We need to consider also the notion that higher level data products can be reconstructed from only Level 0 data. For the data products from NASA's Earth Observing System, many of the investigation teams have contributed very large production code bases that convert the Level 0 data into Level 1, 2, and 3 data. The MISR and CERES teams at ASDC each have about one million lines of code, subdivided into seven to thirteen major subsystems. Furthermore, the code contains or accepts input files containing very large numbers of coefficients that provide instrument calibrations, data editing limits, radiative transfer constants, and numerous other numerical functions. Development of these code bases has involved hundreds to thousands of person-hours. It is reasonable to expect that reconstruction of the higher level data from Level 0 data may require expenditures in excess of \$5M to \$10M for each data set. Reconstructing higher level data products would become very difficult if

more than ten years elapses from the time the data producers have moved on to other projects – even if obsolete hardware could be reassembled. If the data were to be produced on new machines, then it seems highly likely that the validation of the new products would require an investment similar to the original validation budget, based on experience with migrating code from old machines to new ones.

Given the difficulties we have suggested, off-site storage and multiple-copy replication become important aspects of digital preservation of Earth observations. In this case, the features of the Storage Resource Broker that deal with replication become particularly important. Not only does the Storage Resource Broker improve inventory control, but it also speeds replication. The improved replication speed also encourages automated replication, which reduces the cost – and may decrease the probability of error. In this case, grid computing may markedly improve the effectiveness with which we can preserve important data records for such applications as climate change.

OAIS REFERENCE MODEL

Cost is clearly a dominant factor in long-term preservation of data[11]. To the extent that we can create archives that adhere to standard designs and operations models, the more likely we will be able to identify and reduce the factors that increase archival costs. The Open Archive Information System (OAIS) Reference Model (RM) is a very important component in achieving such standardization [12].

The OAIS RM was created by members of the Space Sciences community and has gone through the long, tedious process of being accepted as an international (ISO) standard. It has been well received in the Digital Library community. In outline, the OAIS RM is a functional description of an archive, broken into six distinct areas:

1. Ingest
2. Data Management
3. Data Archival
4. Data Access
5. Administration
6. Preservation Planning

Each of these areas is further broken into a large number of smaller functions, carefully described in the Unified Modeling Language (UML) that makes up the standard. The OAIS RM is thus a useful compilation of the functions an operational archive needs to include.

The challenge for system architects will be to take the abstract description provided by the OAIS RM and map onto it the software architecture needed to operate Linux clusters and Grid computing within the standard. It appears sensible to think that the process of producing such an architecture will start by dealing with what a single archive needs to include – and then later consider how a federated system of archives might develop.

Such development will be particularly interesting if we can use XML to package the interarchive communications and data transfers. The OAIS RM already provides suitable packaging nomenclature in terms of 'Submission Information Packages', 'Archive Information Packages', and 'Dissemination Information Packages'. We note that if the archive and its data are suitably packaged, it does become possible to consider storing the entire archive (including the software and user interfaces) off-line and then reconstructing an active archive from the stored files. Of course, this approach will require careful attention to the information packaging in order to obtain high reliability in the reconstruction.

SUMMARY

In this paper, we have considered several aspects of current data management and how they might be improved through the use of commodity Linux clusters and Grid computing. Current experience is quite hopeful in this regard. We have already achieved proof-of-concept that these elements improve throughput sufficiently that they can substantially improve data access for users. As we have suggested, there is also evidence that the improved capabilities offered by these elements can lead to a new conception of what a data archive may do to increase the value of the archive's contents and the probability of long-term preservation.

This is not to say that the current implementations are perfect. It is clear that Grid computing needs work to simplify its interfaces and to make it easier to install, particularly in the presence of firewalls. However, even with these drawbacks, Grid computing appears headed toward providing a substantial improvement in our ability to provide cost-effective archives.

ACKNOWLEDGEMENTS

The authors are grateful for the support provided by NASA's Earth Science Enterprise. They also appreciate the excellent editing assistance of Ms. Charlene Welch of SAIC International.

REFERENCES

- [1] R. B. Lee, B. R. Barkstrom, G. L. Smith, J. E. Cooper, L. P. Kopia, R. W. Lawrence, S. Thomas, D. K. Pandey, and D. A. H. Crommelynck, "The Clouds and the Earth's Radiant Energy System (CERES) sensors and pre-flight calibration plans", *J. Atmos. and Oceanic Tech.*, vol. 13, pp. 300-313, April 1996.
- [2] G. L. Smith, "Effects of Time Response on point spread function of the scanning radiometer", *Applied Optics*, vol. 33, pp. 7031-7037, October 1994.
- [3] F. O. Huck, N. Halyo, and S. K. Park, "Aliasing and Blurring in 2-D Sampled Imagery", *Applied Optics*, vol. 19, pp. 2174-2181, 1980.
- [4] S. K. Park and Z. Rahman, "Fidelity analysis of sampled imaging systems", *Optical Engineering*, vol. 38, May 1999.
- [5] C. L. Lawson, and R. J. Hanson, *Solving Least Squares Problems*, Prentice-Hall: Englewood Cliffs, NJ, 1974.
- [6] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes: The Art of Scientific Computing*, Cambridge: Cambridge, UK, 1986.
- [7] National Research Council, Committee on Coping with Increasing Demands on Government Data Centers and Committee on Geophysical and Environmental Data, *Government Data Centers: Meeting Increasing Demands*, 2003.
- [8] I. Foster and C. Kesselman, eds., *The Grid: Blueprint for a New Computing Infrastructure*, San Francisco: Morgan Kaufmann, 1999.
- [9] J. Joseph and C. Fellenstein, *Grid Computing*, Upper Saddle River, NJ: Prentice Hall, 2004.
- [10] B. R. Barkstrom, T. H. Hinke, S. Gavali, W. Smith, W. J. Seufzer, C. Hu, and D. E. Corder, "Distributed Generation of NASA Earth Science Data Products", *Journal of Grid Computing* vol. 1, pp. 101-116, 2003.
- [11] B. R. Barkstrom, M. Finch, M. Ferebee, and C. Mackey, "Adapting Digital Libraries to Continual Evolution", *Proc. 2nd ACM/IEEE-CS joint conference on Digital libraries*, July 2002.
- [12] Consultative Committee on Space Data Systems (CCSDS), *Reference Model for an Open Archival Information System (OAIS)*. Blue Book. Issue 1. January 2002. This Recommendation has been adopted as ISO 14721:2003, see http://ssdoo.gsfc.nasa.gov/nost/isoas/ref_model.html.

The ESA Service Support Environment - Exploitation of the long term archives

Sergio D’Elia⁽¹⁾, Pier Giorgio Marchetti⁽²⁾

⁽¹⁾European Space Agency Via Galileo Galilei 00044 Frascati, Italy
email Sergio.Delia@esa.int

⁽²⁾European Space Agency Via Galileo Galilei 00044 Frascati, Italy
email Pier.Giorgio.Marchetti@esa.int

INTRODUCTION

Despite the success in many scientific areas, Earth Observation (EO) has not yet evolved into a real, self-sustaining, service industry. Ongoing efforts to broaden the European EO market highlight the need for EO services and “information products” closer to customer expectations and processes (easily understandable and ready-to-use without further manipulations). In general terms there is the need to move from a “data” or “product” centric vision of EO ground segments to a “service” oriented approach.

Today the lack of an integrated, transparent and service oriented infrastructure adds a further barrier to the transformation from basic EO products (data, images) into information and services. This transformation is performed by a relatively small number of specialised companies operating independently in specific application domains. This separation increases costs and leads to a time consuming process, which prevents the optimisation of allocated resources. Furthermore the exploitation of time series and the access to long term archives which are fundamental for the monitoring of global changes remains cumbersome and practically limited to research organisations.

These problems limit the deployment of cost effective EO services, in particular when support to service orchestration is limited and companies spend their efforts in getting hold of the data or in repeating basic processes outside their own core business, which may be efficiently performed by specialised partners.

Ease and low-cost access to recent and archived Earth Observation data in an environment enabling the setting up of strategic partnerships for providing synergic services are the key factors to the development of an EO service industry. This environment could in turn trigger a further evolution towards a new distributed, flexible, scalable and configurable ground segment permitting orchestration of services for the extraction of information from the wealth of data stored in long-term archives over the last two decades.

In order to cope with the above-mentioned issues, the Agency has established the Oxygen concept with the purpose to:

- Facilitate access to EO data from ESA and other missions
- Increase sustainability of EO data provision, by widening the range of offered data sources
- Make the operations of EO missions more efficient




Figure 1 – Services Environment, with Services Infrastructure and Image Information Mining sites

Along the lines of the Oxygen concept and building on the evidence of the emergence of a Service Environment dimension besides the traditional Data and User Environments, this paper shows how the identification of a set of common EO related standards and the support of a neutral “Services Infrastructure” (see Figure 1), capable to foster the seamless integration of long term archives, EO products and services, may facilitate the set up of open operational services for Earth Observation in Europe.

SERVICE SUPPORT ENVIRONMENT

After 3 years of technological research, ESA proceeds towards a “Service Support Environment” (SSE), which will be the first version of the “Services Infrastructure”. It will support service providers in setting up services, through a new set of interactions necessary mainly to empower the B2B relationships among service providers (see Figure 2). This will facilitate and automate the steps traditionally executed on the basis of personal (human) interactions, which are nice to have in an early definition of the service, but which may become a burden when the service becomes operational and the same tasks have to be repeated for tens or hundreds of time.




Figure 2 – Relations between different user types

The system is open also to service consumers through B2C relations (see Figure 2). They can be of many types, like international organisations, national institutions, commercial users, researchers, service providers, value adders, etc. The SSE (see Figure 3 for the high level architecture) is being implemented with the following objectives in mind:

Permit a very easy service publication, chaining, identification and fruition

It should be possible for a service provider with limited programming skill to publish his services with maximum a few days of work. This is made possible by the TOOLBOX, provided and maintained free of charge by ESA. It permits to easily describe and link the service into the network, with almost no programming knowledge. The TOOLBOX also translates the terms from the service provider legacy environment into the terms used in the SSE.

More important, it shall be possible to chain (or orchestrate) services, via the use of a workflow engine hosted in the SSE Portal. This will permit the creation of complex services from simple ones. Therefore even Service Providers with no services of their own, can create new services just reaching the necessary agreements and chaining existing services.

Services to remain at Service Providers’ sites

It shall not be necessary to move services to a central place or to a service network. They will remain in their current environment under full control of the service provider.

Support for on-line, off-line and subscription services

The environment shall permit the exploitation also of offline services, that is services which require manual intervention and which can also take long time to complete (e.g.: complex interferometric processing, processing based on data still to be acquired, etc.). Subscription shall be permitted for regular, repeated execution of the services.




Figure 3 – SSE high-level architecture

Permit easy, parallel Catalogues' Inter-accessibility

Catalogues shall be accessible via the environment. They shall be seen as services and therefore shall be chainable with other services. This capability has been implemented and tested with EUSC, CNES and Spot Image for inter-accessibility of the SPOT and Envisat catalogues.

Catalogues can be accessed via the spatio-temporal and standard (e.g.: mission, sensor, etc.) attributes as well as via new attributes linked to image content. This has been demonstrated by permitting the identification and selection of MERIS Low Resolution images, free from clouds over the user area of interest.

Based on standards

The most widely used standards from W3C, OGC and others have been selected and used for maximum system flexibility and evolution potential.

Minimum number of commands

The system can work with only two commands: **Request for Quotation (RFQ)** and **Order**. Two additional commands have been added to speed up the catalogue search: **Search**, for parallel catalogue interrogation and results' recombination, and **Present**, for displaying detailed attributes and quick-looks.

Environment "not binding" and "neutrally" managed

Any one shall be free in joining or leaving the environment, which will be managed by an "over partes" entity. The service providers, in the workshop held in 2002, have designed ESA has the organisation providing the highest confidence on neutrality.

Allow services in national language

The network permits the installation of regional portal besides the SSE portal. They can be used to host services in national language. The same services can be repeated in the SSE portal in English for wider penetration within Europe and the world.

Evolve with Service Providers' requirements

An evolution plan has already been identified. It will be focused on a GSTP project for the period 2005 – 2007, which will address new requirements driving from users, services providers, projects (e.g.: the EC Integrated Projects) and programmes (e.g.: GMES, GSE, INSPIRE, etc.).

BENEFITS FOR THE EXPLOITATION OF LONG TERM ARCHIVES

The Agency's experience gained in the Services Infrastructure on service orchestration confirms very large interest from small and large companies. SMEs are particularly interested in exploiting such infrastructure for quick and low cost demonstration and exploitation of new service chains.

This leads in the following future directions for archive services:

1. expand the range of processing which is dynamically performed on the archived products, by allowing the definition of virtual products which are
 - a. either pre-defined by the archive owner
 - b. or defined by the archive user which delegates the archive to execute a "processing" (run an algorithm) he has defined or developed
2. empower new views on the archived products exploiting new approaches (e.g.: from image information mining) which allow the users not only to search product descriptions and metadata (sensor, time, geographical coordinates) but also to select archived images from their content, as well as, more in the future, to directly access the information contained in the images without the need to request and process the images
3. expand the temporal dimension by the definition of multi-temporal archival services which are capable of directly extracting dynamic information from archived image time series
4. expand the sensor dimension by the definition of multi-mission processing / fusion services
5. allow a seamless transition between traditional media (CD-ROM, DVD, Tape) based delivery (and order) services and online (FTP, Web Coverage Server, JPEG2000 clients, etc.) delivery (and order) services

The latter point is of paramount relevance for all long-term archives, which are trying to catch up with new models of customer servicing and product delivery. All above points can exploit the access to physically or virtually co-located archives. The requirements on the archive side would be minimal and mainly related to the adoption of the:

- proposed interface control document [1] based on the widely accepted W3C Web Services standard
- vision of the OAIS model [2] where the consumer is a service itself

RELATED ACTIVITIES AND STANDARDS

A number of additional activities have played a fundamental role in permitting to consolidate the role of the SSE and should be extended to the archives. Among them, of paramount importance is the effort of standardisation that ESA and several national Agencies (CNES, DLR, ASI, ...) are making on the catalogue access, the order and the user visibility over mission planning, which should be extended to include explicitly the archive as "data source".

Furthermore the forthcoming convergence of Web Services and GRID technology confirmed at the Global Grid Forum early in 2004, is paving the way to future additional level of features, which are very promising from an architectural point of view. The most interesting one being a web service resource framework, which removes the notion of persistence from the characteristics of the services, allowing dynamically to create and destroy services on request on a widely distributed set of resources.

RELATIONSHIP WITH OTHER ENVIRONMENTS

The availability of SSE has raised concern on the role that a service infrastructure provided by ESA may play in the context of related environments and services to be made available by programmes like GMES or by initiatives like INSPIRE. From Figures 1 and 2 it is clear that the Service Support Environment, despite being open towards other domains, has been designed from EO specific requirements, with the main emphasis on enabling a natural evolution from a data oriented to a service oriented ground segment for Earth Observation.

Several players will and are mandated to develop regional or thematic service infrastructures. The authors here present for discussion an initial basic scenario, which foresees that service support environments are available not only for EO services, but also for the integration of services based on different data sources like in-situ measurements, meteorological data and in general for services based on geospatial data. The proposed scenario foresees that a Web of Services can initially be very loosely coupled with other services or service support environments dedicated to specific domains (e.g. meteorological services). The advantage of the proposed approach is that no assumptions are made on the architectures owned by other programmes, except that they should be able to interact on the basis of a W3C like “Web of Services” concept, thus allowing service providers and developers the possibility to set up plans for service development and deployment. The disadvantage of this very simple scenario, which relies on a web of very loosely coupled service environments, is that it is intrinsically weaker vs. strong requirements. Therefore complex services, which need multiple data sources and which have strong constraints (e.g. in terms of response time, reliability, resilience, etc. as in services linked to natural events), might not find all the necessary elements for their full satisfaction. However, this very loosely coupled scenario could naturally evolve in future into a more coherent environment, grouping the best experiences made in the various domains.

CONCLUSIONS

The synergy from future activities concerning service support environments and archives should allow fulfilling the objectives to:

- reduce the number of different interfaces (Interface Control Documents) necessary within a ground segment
- permit to design and implement:
 - “virtual archives” composed of elements (e.g. ingestion chain, catalogue, storage) and of services, which can be activated on demand
 - “virtual ground segments”, which may include “virtual archives”
- permit geographic distribution and on demand creation of “virtual ground segments” and “virtual archives” in order to better respond to user requests or to the need to deploy new services of temporary or permanent nature
- increase the flexibility with which new services / products can be provided by individual ground segments or archives
- support both media based delivery and online availability, as well as seamless transition between the two

A flavour of available services from the SSE can be obtained accessing the web site [3]. More specific information on SSE, and on the selection of technologies and standards can be found in [4].

REFERENCES

- [1] “SSE - Interface Control Document” <http://kant.esrin.esa.int/massRef/documentation/icd.pdf>
- [2] “Reference Model for an Open Archival Information System”
<http://ssdoo.gsfc.nasa.gov/nost/wwwclassic/documents/pdf/CCSDS-650.0-B-1.pdf>
- [3] “Service Support Environment Portal” <http://services.eoportal.org>
- [4] Y. Coene, C. Heylen, P.G. Marchetti, “An Environment Supporting Earth Observation and GIS Service Orchestration: The Mass Project”, *IEEE Geoscience and Remote Sensing Newsletter*, December 2003 (Issue 129, ISSN 0161-7869, page 19-22) <http://earth.esa.int/rtd/Articles/IEEEGeoscienceNewsletterDec2003.pdf>

Session Lesson Learnt

Value Added Services for Mars Express' ASPERA-3 Data System

Joey Mukherjee ⁽¹⁾, Michael Muller ⁽¹⁾, Andrew Galus ⁽¹⁾, Sandee Jeffers ⁽¹⁾, Carrie Gonzalez ⁽¹⁾, Jeran Pardue ⁽¹⁾

⁽¹⁾ *Southwest Research Institute
6220 Culebra Road
San Antonio, TX, 78228, USA*

joey@swri.org, mmuller@swri.org, agalus@swri.org, sjeffers@swri.org, cgonzalez@swri.org, jpardue@swri.org

ABSTRACT

Based on the Cluster High Resolution Data System, the Mars Express ASPERA-3 Data System offers more capabilities to the user. First, data is retrieved from the hosting institution within an hour of the data being available. Next, the data is preprocessed to remove packets from: other instruments, duplicate packets, out of sequence packets, etc. The data is then merged with the housekeeping data to ensure the highest quality raw telemetry packet with the data in the proper time sequence with the data in the order needed to process a more scientifically sound telemetry packet. This packet is then processed into a standard data format and then archived into the database. Once archived, the data is then immediately available for analysis via a web-based system or via the independent data analysis system, Southwest Data Display and Analysis System, or SDDAS.

This paper will give the full details of the process from data being transmitted from the satellite to the analysis of the data by the scientist. The key points being made are the value added services of the system. These include the immediacy of the data availability, the integrity of the data, as well as the quality of the plots being made on a near real-time basis.

INTRODUCTION

Southwest Research Institute (SwRI) offers a number of value added services to the data acquisition and visualization aspects of Mars Express' ASPERA-3 instrument. In this paper, seven value added services will be discussed along with our future intentions to further develop our system to add other value added services. There are seven areas which will be the focus of this paper: data availability, data integrity, data plotting, data downloads, data inventory, data access, and data mining. All of these services are available or will be available soon from the web or from Southwest Research Institute's plotting software, SDDAS.

SDDAS, or the Southwest Data Display and Analysis System, is an open source client-server software distribution used for analyzing data and offers many methods for analysis. These include spectrograms, line plots, three-dimensional orbit plots, contour maps in both two and three dimensions, XY plots, and energy angle plots. All three-dimensional plots we create offer the scientist the ability to rotate and zoom the visualization in order to see all of the data from all look directions. Using specially made 3D glasses from Crystal Eyes, the visualization is enhanced even further. The plots will appear to be floating out of the screen in true 3D perspective. In addition to offering the visualizations, SDDAS has a framework by which other groups can add software to SDDAS.

DATA AVAILABILITY

For any system, data availability is critical. Data should be available for use as close to real time as possible. Currently, real time is not possible due to data being retrieved by the European Space Operations Center (ESOC) at seemingly random intervals. Furthermore, once at ESOC, the data is packaged and sent to the principal investigator (PI) institution, which is the Institutet för Rymdfysik (IRF) in the case of ASPERA-3. Currently, our system will poll IRF every hour checking for new data. There is no real reason why this check could not be done more often or less often, depending on the needs of the scientist needing the data and the raw data availability. Theoretically, this differential could approach real time if the needs were there. If new data is found, the data is retrieved, processed, and archived. Once archived, the data is available for immediate use by the public. What this means is data is available to scientists no later than one hour as received by PI institution. Currently, ASPERA-3 data is available only to ASPERA-3 Co-Investigators; however, plots of the data are available to all, but are watermarked before delivery. Technically, there is no reason why all data could not be made available immediately.

DATA INTEGRITY

The data for ASPERA-3 is downloaded from IRF in the ESA Standard Formatted Data Unit (SFDU) format. Unfortunately, raw telemetry is often a mismatch of data packets received from the satellite. Every experiment's data stream is mixed within all the other experiments in potentially non-time-sequential fashion. Furthermore, other problems can also exist in the data such as incomplete packets and duplicate packets. The data must be first separated by experiment and instrument types and duplicate and incomplete packets removed. An additional inconvenience derived from the raw data is that the housekeeping data is delivered separately from the science data. Logically, this makes sense since the housekeeping data is applicable for all experiments; however, without having the housekeeping data, there is no good way to validate the science. To alleviate all these concerns, we have developed pre-processing programs which will rewrite the raw data packet in its native raw format such that all duplicate and bad packets are removed; the data is ordered time-sequential; and the housekeeping data is interleaved with the science data. The end result is a far more functional raw data packet that can be further processed or distributed as needed. Currently, we do not distribute these intermediary files since there have been no requests to that effect. We do, however, use these files to process the data into our own format that is archived for use in SDDAS or via the website. This additional processing results in further integrity checking such that the data delivered is of the highest quality possible.

DATA PLOTTING

Once the data has been processed and archived, it can be plotted in a variety of ways. The plots currently available are spectrograms, line plots, orbital plots, and spectral strip charts. From the web, the user is allowed to specify which plot he wishes to view from a list of predefined plots and set the time. This offers the scientist the quickest way to get to a plot by just selecting one of our many predefined plots and changing the time to a region for which there is data. Once this information is entered, the data is plotted using the *Java Interactive Plotter*. With this, the user is allowed to zoom into a particular time region by marking the boundaries of the region with their mouse or with the click of another button; the user can zoom in using the time delta. To enable the interactive zoom feature, a tracker bar displays the time as the user moves their mouse over the plot. The user is also allowed to page forward and backward through the data using the time delta as the amount of time to move forward or backward. The interface will also keep a history of all the plots the user has made so the user can quickly go to a previously generated plot. If the user wishes to save the plot, he can either save the plot as a GIF or generate a publication quality postscript plot. An example of a plot within the interactive data plotter is shown below.




Fig. 1 Interactive Data Plotter

This interface is only applicable on the spectrograms and line plots. If the user needs more control over the plot, users must use SDDAS that will enable changing virtually any aspect of the plots through the use of the graphical user interfaces (GUI) included. Other advantages of using SDDAS include having full interactivity with the 3D aspects of contour and orbital plots as well as a full suite of software available for retrieving and analyzing data.

Finally, for the most time-efficient method of interfacing with the data, we are also creating summary plots that are generated on a daily basis. These summary plots show the entirety of each data file using a predefined plot layout. Currently, we are only generating summary plots for two of the instruments on the ASPERA-3 instrument. It is our intention to offer summary plots for each of the instruments. Our current summary plots show each of the sixteen anodes of the electron experiment on ASPERA-3 and another shows the sixteen azimuthal angles of the ion mass analyzer compared to two of the anodes of the electron instrument. Both plots also show the spacecraft's position in space in relation to Mars. To view a summary plot, simply pick which time span is of interest. This will be developed more in the future.

DATA DOWNLOADS

Currently for ASPERA-3, data downloading from the web is not yet available, but it is available from SDDAS. However, there is a high probability that web access will be allowed in the near future so it is being presented here for completeness for when it is permissible. Once we have processed the data into our format, it can only be used by software that uses that format. To resolve these issues, we allow the exporting of data from our format into other commonly used formats. The most common format is National Space Science Data Center's Common Data Format, or CDF. Other formats we export into include netCDF (network Common Data Format), XML (extensible Markup Language), ASCII, and a PDS (Planetary Data System) compliant format complete with label files that are directly admissible to the PDS. Eventually from the web, the user will be able to select a data set of interest and a time range. With a click of a button, the data will be converted on the fly to the data format of their choice and be able to be saved to their system. For PDS, there is a requirement that this be automated. Thus, software will be written such that the entirety of the data be converted to PDS format and sent to them. It is conceivable that this be done for any of the formats if another institution wished to archive this data in a different format.

DATA INVENTORY

Being able to see what data is available is almost as important as seeing the data itself. To that end, we have developed two methods of viewing what data is available. The first method is just showing a list of data granules which show a begin time, end time, file size, and the date added to the archive. With this information, one can see how recent the data is should data ever be reprocessed as well as the density of the data should file size be a reasonable method of data density. The unfortunate aspect of this method is that comparing the availability of multiple data sources is difficult. Our most recent method of showing data inventory is our *Interactive Data Inventory*, which is a Java applet that draws a Gantt chart of the data available. The user is allowed to select what data sources and a time range they wish to compare and a bar is drawn representing the data that is available. The biggest advantage of the Gantt method is a quick way to visualize data availability. Additionally, the user can zoom using the same mechanism as on the plots. A tracker bar is overlaid on the chart which the user can highlight a region and zoom showing in more detail data availability. Currently, we can visualize the data availability of granules as small as a nanosecond. Fig. 2 shows a screenshot of the Interactive Data Inventory in action for the Mars Express data and the Mars Global Surveyor (MGS) data from June 1, 2003 to July 15, 2004. The colored areas show areas where we have data. There appears to be continuous coverage for all of MGS for most of the time where we have Mars Express data. Selected data sources for Mars Express include all of the instruments on ASPERA-3. The tree structure allows us to break up the data sources into a natural hierarchy of projects, missions, experiments, and instruments.




Fig. 2 - Example of Interactive Data Inventory showing data availability from June 1, 2003 to July 15, 2004.

DATA ACCESS

Protection of data and seeing who is using the data can be valuable tools in obtaining future grants and acquiring co-investigators. We support a fully open system such that information ought to be freely available. However, we understand that some principal investigators (PI) may not wish to allow others access to the data immediately or ever, or may want to have more control over who gets to see their data. To handle this scenario, we have implemented rights access management into our website. A “root” user is defined and this root user gets to create groups and other administrators. A group is merely a collection of users and an administrator is a user who can add other users to their groups. A user can definitely be a member of more than one group administered by more than one administrator. Every plot or download can be given either anonymous access or individual rights to certain groups. If a plot is labeled as anonymous, even public users can view the plot; otherwise, you must have access to the group to view the plot. For Mars Express, we have two administrators and two groups. The first group is the ASPERA_CoIs who are administered by the PI, Dr. Stas Barabash, who controls who accesses the Mars Express ASPERA-3 data. The second group is the MGS_CoIs who are administered by Dr. Mario Acuña who controls who accesses the Mars Global Surveyor (MGS) data. Dr. Barabash is not allowed to give access to the MGS data and Dr. Acuña is not allowed to give access to the ASPERA-3 data. However, both can add users to the site and see who is using the data. This rights management allows us to add more and more data sets to our site and protect the data for the sake of the PI. Since the PI can give access at their leisure from the web, they can feel safe about giving us their data to process and archive. It is our intention to have all data from Mars Express in our system at some point. Other aspects of rights management are also evident in our plots.

DATA MINING

Data mining is a new up and coming field. We are starting the basics of a data mining engine that could change how scientists look for data. Once the data has been archived, traditionally space physicists have used time as the exclusion factor (i.e. only look at data at this particular time range). However, one must know exactly when an instrument is in a particular mode or region of space. We have developed software that traverse the entirety of our data looking for information. Right now, the data we are looking for are simple instrument mode changes. Technically, there is no reason why we would be limited to just mode changes. We could write a special module that categorizes data by virtually anything including regions in space. Once we find the information we seek, the information is added to a relational database so it can be queried in a relational manner. Eventually, an example of a question we hope to answer is: “show me what data is available from Mars Global Surveyor and Mars Express such that both instruments are over the southern hemisphere of the planet such that no spacecraft is greater than 700km further than 60° latitude.”

TECHNICAL APPROACH / METHODS

SwRI makes extensive use of free, open-source software to develop our software. Within SDDAS, we use Motif [1] (previously commercial, but recently open-sourced by the OpenGroup for all platforms but Microsoft WindowsTM), and Tcl/Tk [2] for our GUIs. Our database is MySQL [3] on the server and our own database for the SDDAS client installs. For visualization, we use the Visualization Toolkit (VTK) [4] and our own graphics package. Our website is run on Linux and uses Apache [5] to serve the web pages. The web pages, including much of the interactive aspects are primarily written in PHP [6] and PERL [7]. Other aspects of our web site are written in Java and we are working to add more Java software within SDDAS. All referenced products are free and open-source making them especially appealing to projects when money is tight.

FUTURE VALUE ADDED SERVICES

In the future, we are hoping to add even more services to our site and SDDAS. These include various mathematical methods such as detrending and moments analysis; and integration of the Interactive Plotter and the Interactive Data Inventory such that both these tools written in Java are coupled such that one can plot a region of time as well as a selection of data sources with a few clicks of the mouse. We are always working on creating more plot types and other interactive tools for use on the web and within SDDAS. Furthermore, we are working on populating our databases with more scientific information to create really fancy queries of data such that we can easily integrate with the SPASE [8], the Space Physics Archive Search Engine. Finally, this could not happen without more data sets to integrate into our system. We hope to have all of Mars Express, as well as other Mars missions, within our system at some time. The more data available, the more elaborate the database that can be queried.

CONCLUSION

We have shown several value added services of our system. In summation: fast automatic data availability makes the act of using the data that much easier for the scientist; high data integrity ensures the scientist quality data; data plotting allows a scientist to view the data with minimal effort; data downloading allows the scientist to retrieve the data in a format of their choice (assuming we can translate to their choice); data inventory lets them see what data is available for given times and instruments; data access allows the PI to protect data or offer data to only people of their choosing; and finally, data mining offers scientists a way to find data without prior knowledge of the data or events.

REFERENCES

- [1] Motif – <http://www.openmotif.org>
- [2] Tcl/Tk – <http://www.tcl.tk>
- [3] MySQL – <http://www.mysql.org>
- [4] VTK – <http://www.vtk.org>
- [5] Apache – <http://httpd.apache.org>
- [6] PHP – <http://www.php.net>
- [7] PERL – <http://www.perl.org>
- [8] SPASE – <http://www.igpp.ucla.edu/spase/index.cfm>

FROM DATA TO KNOWLEDGE IN EARTH SCIENCE, PLANETARY SCIENCE, AND ASTRONOMY

Elaine R. Dobinson, Joseph C. Jacob, and Thomas P. Yunk

*Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 91109-8099
USA*

Email: {Elaine.Dobinson,Joseph.Jacob,Thomas.Yunk}@jpl.nasa.gov

ABSTRACT

This paper examines three NASA science data archive systems from the Earth, planetary, and astronomy domains, and discusses the various efforts underway to provide their science communities with not only better access to their holdings, but also with the services they need to interpret the data and understand their physical meaning. The paper identifies problems common to all three domains and suggests ways that common standards, technologies, and even implementations be leveraged to benefit each other.

INTRODUCTION

The past several decades of scientific exploration of the Earth, the sky, and the entire Solar System from space has produced an abundance of data. Remote sensing and in situ instruments are measuring numerous physical quantities and producing a wide range of numerical and imaging data sets. The curation of these data is, in most cases, the responsibility of the national archives for the appropriate scientific disciplines and subject area, while the analysis of the data is left to the individual science investigator. Thus, within the United States alone, there are the Earth, planetary, and astronomy data archives, each housing terabytes of data collected by satellites flown since the early seventies. The data stored in each of these archives are the original “raw” data, plus various higher-level processed data, usually converted to meaningful physical parameters for scientific study. Access to these data is generally very good, as most digital data are electronically available, and the data archive systems have done a thorough job of “cataloging” their holdings. Yet, the scientific process of converting these data into a real understanding of the universe and its physical phenomena is arduous at best. The sheer volume and diversity of data pose formidable problems to the science investigator, and the data preparation phase of an investigation typically accounts for more than 80% of the task. Furthermore, the same preparation activities – data acquisition, format transformation, re-gridding, re-projection, visualization, etc. – are done repeatedly from investigator to investigator, regardless of the domain, and from one study to the next. However, recent development activities, by information technology (IT) professionals, working together with their science colleagues, have focused their attention to alleviating these problems. In the remaining sections of this paper we examine some of these separate efforts underway in the various science domains listed above, with the goal of recommending common IT solutions across these different domains, and leveraging on emerging IT standards and developments for the benefit of all.

EARTH SCIENCE DATA SYSTEMS

The Earth science data systems have, by far, the largest and most diverse collection of data, having benefited from a wealth of attention on studying our home planet from space. The data resulting from these studies span long time periods, and consist of many remote and in situ measurements, collected by government agencies around the world. While this paper focuses its discussion to only those data collected by NASA and stored in the Earth Observing System Data and Information System (EOSDIS), the problems and solutions discussed have broad application across agencies and across the globe.

Observing the Earth from space entered its prime during the 1980s when the technology of conducting observations from Earth orbiting satellites became mature. The study of global climate change became a world-wide focus and NASA joined other agencies to measure key climate parameters. Archiving these measurements was the responsibility of the individual agencies, and data systems sprung into being at NASA, NOAA, USGS, as well as in Canadian, European, and Japanese space agencies. International committees formed to work through various standards issues so that the data could be shared among all, and these committees are still at work today.

In the 1990's a new era of the Earth Observing System began in the US, and NASA designed and launched multi-sensor platforms that would repeatedly orbit the Earth and measure land, ocean, and atmospheric parameters continuously over a period of many years. Designed as its flagship series of satellites, Terra, Aqua, and the recently launched Aura have sent back hundreds of terabytes of data to help Earth scientists determine the effects of mankind on the Earth system. To house and distribute these data, a new distributed system of data archives was devised, spanning the land, oceans, atmosphere, and other Earth science data sets, and organized in such a way as to hide their physical distribution. Each of the new Distributed Active Archive Centers (or DAACs) was to be part of a larger whole, but was responsive to its subset of the Earth science community. From the beginning, these DAACs were built to a common design and interface, and data format standards were developed to cover the majority of the data. The emerging standard format for these new data (HDF-EOS), based on an existing format, but with geo-location added in, would apply to many of the data collected by the EOS spacecraft, and a series of tools were developed to read and use data stored in this standard format. Thus all data in compliance would be interoperable. This was all well and good, except that not every data producer would use the standard the same way, not all data would be sampled in the same way, not all data would be at the same resolution, and not all data would be collected in the same time period. Thus, additional transformations are still required in order to compare different datasets with each other, not to mention all the legacy data to fit into this "system", leaving the individual scientist with an even more perplexing set of problems to solve. Hence, because these tasks are overwhelming to the individual scientist, each instrument team tends to focus on its own data issues for that instrument – calibration, validation, re-sampling, re-gridding, aggregating – and users of these data tend to work within this team's purview. Interdisciplinary investigations, that is, investigations that integrate data from more than one instrument and more than one topic area, remain extremely difficult. One notable exception to this is the emergence of widely accepted standards for geospatial data, promulgated by the OpenGIS Consortium, which has made tremendous inroads into this problem for all geo-spatial data. However, generic solutions for other data types and other domains still are lacking.

Several recent efforts are currently underway at NASA to aid the interdisciplinary investigator in this dilemma, and to promote the study of Earth System Science and make it feasible. One such effort begun this year is the General Earth Science Investigation Suite (GENESIS) project at the Jet Propulsion Laboratory (JPL) [Yun04]. GENESIS is a NASA-sponsored partnership between the Jet Propulsion Laboratory, academia, and three NASA data centers to develop a new suite of web services tools to facilitate multi-sensor investigations in Earth System Science. These tools will offer versatile operators for data access, subsetting, registration, fusion, compression, and advanced statistical analysis. They will first be deployed in a model server at JPL, and later released as an open-source toolkit to encourage enhancement by independent developers. While the initial work will focus on four premier atmospheric sensors – AIRS, MODIS, MISR, and GPS – the modular design offers ready extension and reuse on many Earth science data sets.

At the core of GENESIS is its scientific workflow engine known as *SciFlo*. SciFlo combines four core ideas to promote software reuse and create a marketplace for science analysis services: loosely-coupled distributed computing using SOAP; exposing scientific analysis operators as SOAP web services; specifying a data processing stream as an XML document; and a dataflow execution engine for a parallel execution plan and load balancing. The SciFlo design grew out of the pressing needs of scientists active in studies with these new sensors. The tools themselves will be co-developed by atmospheric scientists and information technologists from several institutions. At each step the tools will be tested under fire within active investigations, including cross-comparison of spaceborne climate sensors; cloud spectral analysis; upper troposphere-stratosphere water transport; and global climate model testing. The GENESIS tools, eventually to be inserted into routine DAAC operations, will help to inaugurate Earth System Science and will advance a modern data system architecture for realizing the broader vision of NASA's Earth Science Enterprise.

PLANETARY DATA SYSTEMS

The oldest of the NASA data archive systems still in existence today, the Planetary Data System (PDS), was begun in the mid-eighties, as an early prototype of the distributed archives model. Based on recommendations by the National

Academy of Sciences, NASA built an archive system of distributed discipline data centers (called nodes) under central management and development. These PDS nodes conceptually remain in existence today, and are responsible for the curation of the entire suite of data collected by all US and non-US planetary missions, dating back to the early days of the Viking and Voyager missions. Until the advent of the recent Mars orbiting missions, however, the entire planetary science data collection remained small (less than 3 terabytes), and access was easily provided on CD-ROM. However, with the increased number of orbiting missions to Mars, the Galileo and Cassini orbiting missions to Jupiter and Saturn, and the increased resolution of the instrument measurements, the size of the PDS archive now exceeds 300 terabytes. What was once a manageable problem suddenly grew to Earth data system proportions, and the PDS must now provide the same tools and access mechanisms that their Earth science colleagues have come to take for granted. Data must be accessible on-line as collections are too large to span a CD-ROM or two. The PDS has recently developed a new, distributed, on-line access to its data, seamlessly integrating all data from all nodes from all missions. This distributed system resembles the EOSDIS system discussed above.

PDS from the beginning developed a standardized set of “objects”, and a standard nomenclature for referring to these objects. They also developed a standard language for describing the binary files containing these objects. This level of standardization works well within the system, but interoperability outside the PDS, with other widely accepted standards, is poor. This limits the software that can be used to manipulate the data, and the task of converting PDS objects into standard formats is left to the user. While PDS itself attempts to provide some rudimentary display tools, there are significant problems in integrating data across missions and across instruments. For, while the data object language is standard, the coordinate systems and projections used for the data are not. Thus the data sets, even from the various Mars missions, must first be transformed into a common system before they can be used together.

A pending effort by NASA is proposing to solve this problem for the terrestrial planets by leveraging off the Earth science work in the geo-spatial world. Extending the OpenGIS model to include planets other than Earth, and extending the standard interface protocols (WMS and WCS) to apply to the Moon and Mars is part of an effort underway at JPL. The end result will be a server of standardized “maps” of all terrestrial planets, usable by all OpenGIS software. The task of co-registering the individual images, and of re-projecting them to a common standard will still need to be done, but once done, these data will be available for all to use.

ASTRONOMY DATA SYSTEMS

Astronomical datasets exist in the form of images, catalogs, spectra, time series, and numerical simulations, which embody our knowledge of the universe at wavelengths that span the electromagnetic spectrum, from radio waves through gamma rays. Rapid technology improvements in detectors, telescopes, computing, communications, and storage, have given rise to all-sky surveys and precipitated exponential growth in the size of these datasets [Bru02]. As an example, in 1983 the InfraRed Astronomical Satellite (IRAS) captured nearly the entire infrared sky in four wavelengths to yield less than 1 GB of imagery. Contrast this with more recent surveys such as the Digitized Palomar Observatory Sky Survey (DPOSS) and Sloan Digital Sky Survey (SDSS) in the visible wavelengths and the Two Micron All Sky Survey (2MASS) in the near-infrared wavelengths, which each host data at the multi-terabyte level, a 4 orders of magnitude leap in data size in little over a decade. Proposed missions like the Large Synoptic Survey Telescope (LSST) will continue this progression to the petabyte level.

The international astronomy community widely recognized that in the 1990's a significant and ever growing gap had emerged between the size and complexity of the datasets being captured and our ability to effectively extract the wealth of information inherent in them, in order to maximize the science impact of our space missions. This is an information technology challenge that is exasperated by the fact that the data are distributed and served via disparate mechanisms, and also a sociological challenge in that community standards are needed for data, query and search mechanisms, and computational services. A number of “Virtual Observatory (VO)” projects are addressing this problem, including the National Virtual Observatory (NVO) in the USA [Djo02], AstroGrid in the UK, and others elsewhere in Europe and Asia. The International Virtual Observatory Alliance (IVOA) is an effort to coordinate the various VO projects and encourage cooperation on issues that they have in common, such as data formats and interfaces between services.

The Flexible Image Transport System (FITS) has long been a lingua franca among the international astronomy community for sharing astronomical imagery. The FITS format encapsulates the image data with a header containing keyword-value pairs that describe the image and how pixels map to the sky. This meta-data information specifies image dimensions, pixel sampling, location on the sky of a specific reference pixel in the image, coordinate system, projection, arbitrary rotations, and others.

A common operation in astronomy is to extract a catalog of celestial objects from an image. An XML format for catalogs called VOTable has been proposed by the NVO, supporting hierarchical representation of metadata using custom tags [Wil02]. One effort to standardize the form of these tags is the Unified Content Descriptor (UCD), a proposal by the IVOA to define a formal vocabulary to express unambiguously the semantics of astronomical concepts in a human- and machine-readable form [Der03].

The NVO has also defined a set of standards for data providers that define how images and catalogs may be searched and served to the community. The Cone Search protocol is a simple mechanism that enables catalog searches based on a simple center location on the sky and a radius. A VOTable is returned containing the objects that fall within the specified bounds. The Simple Image Access Prototype (SIAP) specification is an extension of this search capability for images instead of catalog objects.

One key theme in the virtual observatory is the idea that new technology and standards will enable easy access to multiple datasets that can be used jointly for multi-wavelength science. The science drivers for this include: (i) search for and study of brown dwarfs, which are faint in the visible wavelengths, but may show up in the infrared; (ii) study of pulsars, which radiate strongly, but with dramatically different characteristics, at a wide range of wavelengths from radio through gamma ray; (iii) study of quasars, which may be distinguished from stars by their unique spectral signature across multiple wavelengths; (iv) search for new objects that are so faint as to be indistinguishable from noise in a single wavelength, but may be identified by correlation across wavelengths; (v) identifying similar classes of objects through clustering in multi-wavelength space; and (vi) discovery of entirely new types of objects with unusual spectra by examination of those outlying objects that do not cluster well in this multi-wavelength space [Djo01].

The technical challenges to be overcome in data federation are in how to relate objects in one archive to objects in another, how to handle the situation where objects appear in one archive but not in another, and how to provide uniform access to distributed data archives hosted by different organizations. This entails providing a data federation layer that hides the underlying search and access mechanisms of the multiple distributed data archives. Furthermore, data federation enables queries and searches that span multiple archives.

Since instruments can provide either high spatial resolution or wide area, but not both, an all sky survey typically archives many thousands or even millions of images, each with a different projection corresponding to the pointing of the telescope. Image mosaics are essential in order to enable the study of objects that span multiple images or for study of star formation regions or large-scale structure of the universe. Construction of an image mosaic entails reprojection of the input images to the output coordinate frame, matching the background intensities across the images, and combining the images to produce a single output mosaic. Several image mosaicking projects exist, including JPL's yourSky custom mosaic web portal [Jac02a], a follow-on project at Caltech and JPL called Montage for flux preserving mosaics [Berr02], and SWarp from the French TERAPIX center [Bert03].

The virtual observatory community has widely recognized that grid computing is a natural fit for astronomy because the data, compute resources, and domain expertise are all distributed. The Grist project led by Caltech is architecting a framework for interoperable services for astronomy compliant with NVO, grid, and web services standards. A number of services will be deployed in support of astronomical data mining, including services for data access, mosaicking, extracting source catalogs from images, clustering, catalog manipulation, statistics, and visualization. An interactive workflow system will allow a scientist to use a visual programming interface to link together these distributed services as needed and control service deployment and execution from a desktop computer.

Computational grids are being used for image mosaicking, which is both a compute- and data-intensive operation. The yourSky portal has been extended into yourSkyG, a mosaic service on the Information Power Grid (IPG), NASA's computational grid infrastructure. In addition, the Montage mosaic software has been deployed as a service on the TeraGrid, NSF's 20-teraflop computational grid [Berr04].

A number of standard visualization tools are widely used and suitable for modest size data analysis, including SAOImage DS9 [Joye03], and OASIS [Good03]. For larger datasets, two applications developed at JPL enable large-scale visualization of astronomical images and catalogs using high-end or low-end resources. The Electronic Light Table (ELT) provides high-performance visualization driven by supercomputers connected to multi-screen "Powerwall" displays [Jac00]. Visualization of selected datasets on standard desktop computers is provided by a web-based system called SkyLite [Jac02b]. Key features of these visualization tools include interactive, high-performance pan and zoom

on multi-wavelength and multi-resolution datasets that are larger than typical memory sizes, and the inter-relation of numerous catalog and image layers to each other.

UNDERLYING THEMES AND RECOMMENDATIONS

We can see from the discussions above that, no matter what the science domain or the subject matter, there are common themes that run through the data systems. An analysis of these undercurrents may give us some clues as to how currently separate efforts may be able to share development resources and leverage common solutions. We have seen the benefit of standards *within* a single discipline; we have yet to realize the benefit of standards *across* disciplines. If we consider these from an information technology perspective only, noting the areas where discipline differences may warrant special cases, we may be able to bring about a greater level of cooperation, and therefore, *a higher level of progress*, than before.

Data Access and Utilization

All of the archive systems need to provide better access to data given the increasing volumes that are now on hand to sort through. Faster searches, narrower selections, and speedier transmission of resulting information are on everyone's list of improvements to be made. The relatively "raw" data stored in the archives must be processed to additional constraints (subsetting, re-gridding, re-projecting, etc.) that cannot be apriori specified, as the constraints depend on the particular investigation at hand. However, the notion of the scientific workflow system to execute this string of processing constraints on combinations of data sets and parameters is common throughout. While the individual operators may have to be data set specific, or at least data format specific, the environment for running them can be generic. Web services and grid services facilitate remote computing and distributed access to data. These services are all part of the new efforts underway, and the use of emerging standards for interfacing to these services can lead to far greater cooperation than before. Joint registries of specific grid services that appeal across science domain boundaries would facilitate publishing and discovery of these common services.

The barriers to this level of cooperation are more social than technical in that the difficulty lies in the communication between what tend to be isolated development teams. Cross team discussions are often difficult, but more and more of that is happening at conferences and meetings such as this one. The more this takes place the more leveraging of work can occur. The emergence of the grid-computing paradigm, originally for high-energy physics, has fostered this level of interchange, and discussions of grid computing is catching on and spreading to all the science domains. While not all applications warrant using a high-performance computing grid, the adherence to the grid data standards facilitates its use where it is advantageous.

Intelligent Archives and Data Warehouses

Another paradigm emerging in the Earth science realm is that of the "intelligent archive" [Ramapryian et al]. In this paradigm, the access and utilization tools reside *within* the data archive system and are brought into execution as needed to satisfy a processing request on demand. The archive's "intelligence" is in determining what needs to be invoked without being specifically told. The scientist specifies what he needs; the intelligent archive figures out how to give it to him. The GENESIS system, when resident at a DAAC, is a flavor of this concept. Format translations, subsetting algorithms, higher-level parameter processing, data mining algorithms, all run on the archived data to produce more meaningful data for the scientist, are all examples of this approach. A related concept in astronomy is the "virtual data" idea being explored in projects such as Grist, which refers to intelligent caching mechanisms in a workflow that enable data products to be drawn from a pre-computed collection when available, and dynamically computed only when necessary.

These "on-the-fly" transformations are in contrast to stored "warehouses" of pre-processed data. This concept also has its place in the scientific world. Large stores of data are served out in the geo-spatial world by Web Mapping Servers (WMS) to promote interoperability and layering of widely dispersed data sets. Indeed, the entire field of Geographic Information Systems (GIS) has been revolutionized by the availability of these standard data sets. Raster images from remote sensing instruments and vector data of known geologic features can be overlaid to produce geologic maps.

Earth scientists have long known the value of these standards; the planetary science community is recently realizing its benefits. The advent of planetary rovers has increased the importance of this technology and the use of GIS for terrestrial planets will greatly enhance these missions. In astronomy, relating multiple wavelength images to each other and to catalog objects is a key theme in a number of data analysis and visualization systems, analogous to GIS in the Earth and planetary domains.

Virtual Observatories and Data Federations

Although the term “virtual observatory” originated in the astronomical community, the concept is widespread in the Earth and planetary realms. Put simply, a virtual observatory (VO) is the amassing of data collected and archived by individual groups to yield new results in combination with other such collections. This “federation” of data thus becomes a virtual observatory for the scientist. Unions of data collections span organizations and countries; the secret to interoperability among them lies in the agreed upon standards used to represent the data and used to interface the data to common computing services. Thus there is arising a federation of earthquake data in the Solid Earth Virtual Observatory (SERVO) based on the astronomy example. Virtual Mars, a similar concept in the Planetary Data System of integrating all of the known data about Mars into standard representation, is another application.

Federated Earth science data systems are perhaps the most mature of the data federations. The concept was first prototyped with the Earth System Information Partners, a union of the DAACs and competitively selected Earth science applications to promulgate the data collected by the NASA Earth science missions and render them more amenable to commercial use. The federation still exists on its own, and NASA has recently added a new round of information partners selected in the REASoN awards to further promote research with and application of Earth science data.

CONCLUSION

This paper has examined three NASA science data archive systems, from the Earth, planetary, and astronomy domains respectively, from the point of view of transforming data into knowledge. We have discussed some of the various efforts underway to provide the science communities with not only better access to the data holdings, but also to provide them with the services and tools they need to interpret the data and better understand their physical meaning. We have identified selected themes common to all three domains and suggested ways that common standards, technologies, and even implementations could be leveraged to benefit each. It remains for us as developers of these systems to ensure that we work together to promote this cooperation.

ACKNOWLEDGEMENT

Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

- [Berr02] G. B. Berriman, D. Curkendall, J. Good, J. Jacob, D. S. Katz, T. Prince, and R. Williams, *Montage: An On-Demand Image Mosaic Service for the NVO*, Astronomical Data Analysis Software and Systems (ADASS) XII, October 2002, Astronomical Society of the Pacific Conference Series, eds. H. Payne, R. Jedrzejewski, and R. Hook.
- [Berr04] G. B. Berriman, E. Deelman, J. Good, J. Jacob, D. S. Katz, C. Kesselman, A. Laity, T. A. Prince, G. Singh, and M. H. Su, *Montage: A Grid Enabled Engine for Delivering Custom Science-Grade Mosaics on Demand*, Proceedings of SPIE Vol. 5493, SPIE Astronomical Telescopes and Instrumentation: Optimizing Scientific Return for Astronomy Through Information Technologies Conference, June 2004.
- [Bert03] E. Bertin, *SWarp v2.0 User's Guide*, 2003.
- [Bru02] R. J. Brunner, S. G. Djorgovski, T. A. Prince, and A. S. Szalay, Massive Datasets in Astronomy, in Handbook of Massive Data Sets, Kluwer Academic Publishers, 2002, pp. 931-979.
- [Der03] S. Derriere, F. Ochsenbein, T. Boch, and G. T. Rixon, *Metadata for the VO: The Case of UCDs*, Astronomical Data Analysis Software and Systems XII, 2003.
- [Djo01] S. G. Djorgovski, A. Mahabal, R. Brunner, R. Williams, R. Granat, D. Curkendall, J. Jacob, and P. Stolorz, *Exploration of Parameter Spaces in a Virtual Observatory*, Proc. SPIE Vol. 4477, p. 43-52, Astronomical Data Analysis, Eds. Jean-Luc Starck and Fionn D. Murtagh, November 2001.
- [Djo02] S. G. Djorgovski, et al., National Virtual Observatory Science Definition Team Report: *Towards the National Virtual Observatory*, April 2002.
- [Good03] J. C. Good, M. Kong, and G. B. Berriman, *OASIS: A Data Fusion System Optimized for Access to Distributed Archives*, Astronomical Analysis Software and Systems XII, ASP Conference Series, Vol. 295, eds. H. E. Payne, R. I. Jedrzejewski, and R. N. Hook, 2003.
- [Jac00] J. C. Jacob and L. E. Husman, *Large Scale Visualization of Digital Sky Surveys*, Virtual Observatories of the Future Conference, June 2000, Astronomical Society of the Pacific Conference Series, Volume 225, 2001, eds. R. J. Brunner, S. G. Djorgovski and A. S. Szalay, pp. 291-296.
- [Jac02a] J. Jacob, R. Brunner, D. Curkendall, G. Djorgovski, J. Good, L. Husman, G. Kremenek, and A. Mahabal, *yourSky: Rapid Desktop Access to Custom Sky Image Mosaics*, SPIE Astronomical Telescopes and Instrumentation: Virtual Observatories Conference, August 2002.
- [Jac02b] J. C. Jacob, G. Block, and D. W. Curkendall, *Architecture for All-Sky Browsing of Astronomical Datasets*, Astronomical Data Analysis Software and Systems (ADASS) XII, October 2002, Astronomical Society of the Pacific Conference Series, eds. H. Payne, R. Jedrzejewski, and R. Hook.
- [Joye03] W. A. Joye and E. Mandel, *New Features of SAOImage DS9*, Astronomical Analysis Software and Systems XII, ASP Conference Series, Vol. 295, eds. H. E. Payne, R. I. Jedrzejewski, and R. N. Hook, 2003.
- [Ram02] H. K. Ramapriyan, S. Kempler, C. Lynnes, G. McConaughy, K. McDonald, R. Kiang, S. Calvo, R. L. Roelofs, D. Sun, *Conceptual Study of Intelligent Data Archives of the Future*, 10th NASA Goddard Conference on Mass Storage Systems and Technologies and 19th IEEE Symposium on Mass Storage Systems, April 2002.
- [Wil02] R. Williams, F. Ochsenbein, C. Davenhall, D. Durand, P. Fernique, D. Giarretta, R. Hanisch, T. McGlynn, A. Szalay, and A. Wicenec, *VOTable: A Proposed XML Format for Astronomical Tables*, 2002.
- [Yun04] T. Yonck, B. Wilson, A. Braverman, E. Dobinson and E. Fetzer, *GENESIS: The General Earth Science Investigation Suite*, NASA Earth Science Technology Conference, July 2002.

U.S. GEOLOGICAL SURVEY SCIENTIFIC RECORDS APPRAISAL TOOL

PV-2004, “Ensuring the Long-Term Preservation and adding Value to the Scientific and Technical Data”

5-7 October 2004, ESA/ESRIN. Frascati, Italy

J. L. Faundeen⁽¹⁾, C. J. Wippich⁽²⁾

*⁽¹⁾ U.S. Geological Survey
EROS Data Center
Mundt Federal Building
Sioux Falls, SD 57198 USA
faundeen@usgs.gov*

*⁽²⁾ U.S. Geological Survey
Geographic Information Office
12201 Sunrise Valley Drive
Reston, VA 20192
cwippich@usgs.gov*

ABSTRACT

The U.S. Geological Survey (USGS) is an agency that gathers, analyzes, archives, and distributes scientific information related to our environment. The USGS uses this information to investigate physical and biological processes and how they are affected by human activity throughout the landscape.

Timely, high-quality scientific information is critical to the decision-making processes that form responsible resource management. This scientific information, maintained as a U.S. Federal record, provides the foundation for future scientists. Beyond legal requirements, the science of records management ensures that these records are preserved and accessible.

Until recently, USGS called for retaining all records without fully evaluating their long-term scientific value. The sheer volume and increased pace of creation of scientific information compels us to evaluate these records in a more systematic fashion and dictates we husband our resources to retain and preserve the most scientifically valuable records.

Appraisal of scientific information is an important tenet of good records management, but it is also the most subjective. Predicting what scientific information may have the most value in the future is a formidable challenge. To address this problem, the USGS has developed an electronic, web-based appraisal tool based upon extensive literature research that

incorporates applicable national and international best practices and standards. This tool uses sequential questions to help appraise and assess the long-term scientific value of collections.

BACKGROUND

The USGS serves as an independent, fact-finding agency that collects, monitors, analyzes, and provides scientific understanding about natural resource conditions, issues, and problems. Because it has no regulatory or management mandate, the USGS provides impartial science that serves the needs of our changing world. The diversity of scientific expertise enables the USGS to carry out large-scale, multi-disciplinary investigations that build the base of knowledge about the Earth. The USGS focuses its scientific investigations into four major areas: natural hazards, natural resources, information and data management, and the environment. These areas are also found in the global community, presenting unprecedented opportunities for the science of the USGS to make substantive and life-enhancing contributions for the United States and the world.

RECORDS MANAGEMENT

For more than 50 years the U.S. Federal Records Act has required agencies to create and maintain adequate documentation of their policies and official business transactions [1]. Complete and accurate scientific and administrative records are essential components of the USGS. Our collective challenge at the USGS is to make sure that the records we create or receive are properly managed through the full life cycle of creation, maintenance and use, and disposition. Ensuring proper records management allows scientists and decision-makers of today and those not born yet access to the substantial science records this agency produces.

A scientific record is any recorded information relating to the work of the agency, regardless of who created it or how the information was recorded. While new technology has made the creation of records easier, many records are paper documents such as letters, completed forms, memorandums, directives, photographs, maps, engineering drawings, and reports. Scientific records increasingly appear in electronic and other forms such as microfilm, motion pictures, sound records, computer tapes, hard or soft disks, and three-dimensional objects such as core and soil samples. Non-record material generally is composed of convenience copies, library holdings, and stocks of publications.

Records management at the USGS involves viewing scientific records through their lifecycle consisting of (a) their creation or acquisition; (b) the use, maintenance and preservation of the records; and (c) final disposition when the records have outlived their usefulness or relevancy to the creating agency. Records management typically involves the processes of appraisal, accession, arrangement, description, access, reference, preservation, disposition, outreach, and advocacy. The appraisal function helps determine what new collections should be accessioned into an archive. Re-appraisals will determine if existing holdings should continue to be retained.

APPRAISAL PROCESS

The 1950 U.S. Federal Records Act dictates that U.S. Federal records having “sufficient historical or other value” warranting their continued preservation become the main objective for appraisers, i.e. determining what collections justify spending resources to preserve them and make them accessible for researchers today and in the future [2]. This deceptively simple process is often one of the most difficult for archivists and records managers. By its very nature, appraisal is the most subjective element within records management. Judgments are passed, many of which have long-term impacts.

USGS TOOL

The USGS Earth observation archives are diverse, containing film obtained from cameras flying in former 1960s intelligence satellites and hand-held photography obtained from NASA Shuttle astronauts to digital images generated from satellites currently orbiting. The scientific collections came to the USGS for many reasons, some of which were never fully documented. With continuing tight budget environments and organizational missions evolving, the USGS began to reappraise its holdings to determine if they are still relevant to the agency's mission. A means to objectively review the collections was needed. Extensive searching in appraisal documentation lead to the development of a web-based USGS Scientific Records Appraisal Tool that attempts to (a) ask all of the pertinent questions about a collection necessary to make an informed decision and (b) document the information assembled to help form a decision to accept a new collection or retain old collections. While the actual questions are critical to the usefulness of the tool, it was also important to compile a history of what was involved in the overall decision-making process. What has evolved can be viewed as an "appraisal checklist" of criteria by which collections can be measured [3].

Once assembled and reviewed repeatedly by USGS records managers and U.S. National Archives and Records Administration staff, the tool was added to the internal USGS web site. A mock appraisal of an existing collection, the Apollo hand-held photography, was conducted. Fig. 1 illustrates a frame from the Apollo collection [4]. The tool was modified based upon the mock exercise. A few weeks later, the same collection was reviewed as the first official use of the tool. Some minor revisions have been incorporated subsequent to the first use of the appraisal and it is now considered to be operational.




Fig. 1. April 1972 Apollo image of the Earth

In assembling the questions that make up the appraisal tool, several sources were consulted including the Society of American Archivists “Selecting and Appraising Archives and Manuscripts” [5] and the recently released International Organization for Standardization (ISO) Records Management standard [6]. Sections were developed addressing mission relevancy, policy, cross-cutting general areas, physical media, metadata, and cost/benefit factors. The appraiser submission part and the mission relevancy portion are listed in Fig. 2. If many of the choices are recorded as “NO” in this area, serious doubts arise about accepting or retaining a collection. Assistance is provided to the appraiser through web links to additional documents such as mission statements, the relevant collection policy, and the pertinent ISO records management elements.



Records Appraisal Tool

This tool assists the USGS in appraising record collections that are offered to, or sought by, the USGS.

Date:

Appraiser:

Collection Name:

Email:

Mission Relevancy

Do the records fit within the scope of our Collecting Policy?	YES ☐	NO ☐	N/A ☐
Does the anticipated current and future utility of the data fit within the EDC mission?	YES ☐	NO ☐	N/A ☐
Are the records significant or unique to the remote sensing, cartographic, and Earth science data user community?	YES ☐	NO ☐	N/A ☐

Fig. 2. Appraiser submission and mission relevancy portions

The policy portion, to a large degree, relies upon the ISO Information and Documentation-Records Management standard. Web links are provided explaining the ISO precepts of authenticity, reliability, integrity, and usability. As in the mission relevancy section, if numerous “NO” answers are recorded, the viability of accepting or maintaining a scientific collection becomes questionable. Fig. 3 illustrates this portion of the tool.

Policy

Can the records <u>Authenticity</u> be judged, i.e. are the records considered to be authentic? (ISO 15489-1:2001(E))	YES ☐	NO ☐	N/A ☐
Can the records <u>Reliability</u> be assessed? (ISO 15489-1:2001(E))	YES ☐	NO ☐	N/A ☐
Are the records <u>Integrity</u> of a high nature? (ISO 15489-1:2001(E))	YES ☐	NO ☐	N/A ☐
Are the records <u>Usability</u> conducive to our anticipated exploitation of the information value in the records? (ISO 15489-1:2001(E))	YES ☐	NO ☐	N/A ☐

Fig. 3. Policy section

The cross-cutting general section addresses facets not easily contained in the other portions. Questions deal with who created the records, i.e. where does the provenance rest, characteristics of the collection such as temporal and spatial elements, and what format the records are in. A sub-set of this portion is represented in Fig. 4.

General Section			
Who created the record and for what purpose?			
What significant contributions, such as unique or under-recorded spatial or temporal coverage, does the collection bring to EDC as defined through the archive's Collection Policy?			
What is the spatial area covered by the collection?			
What is the temporal range(s) the collection spans?			
How does the record meet the information needs and interest of various user groups served by the archive?			
What is the potential utility of the record based on past and present research use?			
What are the physical, intellectual, or legal barriers in making the records accessible?			
Do the records represent a complete population or universe, or a statistically valid sample?	YES ☐	NO ☐	N/A ☐

Fig. 4. General section of the appraisal tool

Typically, the physical media that the records reside on are significant barriers themselves due to hardware, software, and media obsolescence. Figure 5 represents the questions from this section that can greatly characterize the level of effort required to accept or maintain a scientific collection.

Physical Section			
Are the records provided at the level of processing that best preserves the integrity of the records and is the most useful to the anticipated requestors of the data?	YES ☐	NO ☐	N/A ☐
Are the data uncompressed?	YES ☐	NO ☐	N/A ☐
Do the records reside on media that is compatible with USGS systems and have at least 5 years of reliable life remaining?	YES ☐	NO ☐	N/A ☐
Is the format that the records are to be transferred in non-proprietary and computer-compatible?	YES ☐	NO ☐	N/A ☐
If the format is proprietary, has a formal sunset date when the data would be considered Public Domain agreed upon?	YES ☐	NO ☐	N/A ☐
Can the record's file naming convention be documented?	YES ☐	NO ☐	N/A ☐
Are non-proprietary raw, raster-formatted browse available for each record?	YES ☐	NO ☐	N/A ☐
Are the initial, and any subsequent, processing histories available?	YES ☐	NO ☐	N/A ☐
Is the medium that the records reside on stable in its current form?	YES ☐	NO ☐	N/A ☐
Are the records at the optimal generation level for long-term preservation?	YES ☐	NO ☐	N/A ☐

Fig. 5. Physical media questions

Metadata has become one of the most important and challenging elements for long-term preservation and accessibility of scientific collections. As such, the completeness or inadequacy of a collection’s metadata must be completely understood to best determine the levels of service an archive may be able to provide for the collection and in what timeframe that service can be offered. Fig. 6 details metadata questions provided by the tool.

Metadata Section			
Is complete metadata, capable of supporting Federal Geographic Data Committee/International Standards Organization collection- and record-level standards with emphasis upon complete frame/image center and corner latitude and longitude coordinates, available?	YES ☐	NO ☐	N/A ☐
Can the metadata be provided via flat, American Standard Code for Information Interchange, delimited files and indexed to tie to the physical inventory of records?	YES ☐	NO ☐	N/A ☐
Is additional information about the records available? Examples include libraries of documentation, guides, Data Information Files, fact sheets, Frequently Asked Questions sheets, instrument documentation, Preliminary Design Reviews, Critical Design Reviews, lessons learned, hardware documentation, firmware documentation, engineering models, computer models, platform documentation, algorithm documentation, URLs, Principle Investigator contact, ATBDs. 	YES ☐	NO ☐	N/A ☐
Would training be available from the records provider or creator?	YES ☐	NO ☐	N/A ☐

Fig. 6. Metadata questions

Most appraisers would prefer to not have to make decisions based upon cost. It is, however, a constant factor in most appraisals simply because reality dictates that not every scientific collection can be managed, preserved, and made accessible to the level that many would desire. Fig. 7 attempts to capture this difficult, but important aspect of the decision-making process.

Cost / Benefit			
Are there sufficient funds to acquire, maintain, and make records available now and in the future?	YES ☐	NO ☐	N/A ☐
Is there a potential of cost sharing for capital investment and/or recurring costs?	YES ☐	NO ☐	N/A ☐
Would the data be difficult or expensive to replicate by us or someone else?	YES ☐	NO ☐	N/A ☐
Are there significant costs or consequences to the program or the Government if the data are not obtained or maintained?	YES ☐	NO ☐	N/A ☐
Does the estimated research value of the data exceed the costs to maintain them for secondary use by Government researchers or other?	YES ☐	NO ☐	N/A ☐
Estimate the cost of preserving the record weighed against the benefit of retaining the information, i.e. what are the costs of identifying, appraising, and accessioning the records?			
What are the costs of processing the collection to an accessible level?			
Are the resources necessary for any preservation or access functions available?	YES ☐	NO ☐	N/A ☐
What are the annual costs of housing the original records or reducing their bulk by sampling?			
Will special equipment be required to read or process the records?	YES ☐	NO ☐	N/A ☐
Estimate the cost to deaccession/purge the collection. 	N/A ☐		

Fig. 7. Cost/benefit

SUMMARY

Appraising scientific collections is subjective by nature but, by assembling the relevant questions and capturing the answers to those questions, decisions made can be explained and justified, even years later. The USGS Scientific Records Appraisal Tool is still in its infancy and will undoubtedly evolve as it is put into practice more. It is but a means for appraisers to draw from when faced with the challenging task of reviewing scientific collections. The tool is now publicly available via URL <http://edc2.usgs.gov/government/RAT/tool.asp> and comments would be appreciated by the USGS (faundeen@usgs.gov).

REFERENCES

- [1] United States Federal Records Act. Washington. Title 44 United States Code chapters 21, 29, 31, 33 (1950). URL <http://www.ed.gov/policy/gen/leg/fra.html>
- [2] Daniels, Maygene F. and Walch, Timothy, ed. *A Modern Archives Reader: Basic Readings on Archival Theory and Practice*. Washington: National Archives Trust Fund Board, 1984, p. 88.
- [3] Ellis, Judith, ed. *Keeping Archives*. Port Melbourne: D W Thorpe, 1993, p. 172.
- [4] U.S. Geological Survey. Apollo documentation. 2004. URL <http://edc.usgs.gov/archive/nslrda/background/Apollo.html>
- [5] Ham, F.G., *Selecting and Appraising Archives and Manuscripts*. Chicago: Society of American Archivists, 1993.
- [6] International Organization for Standardization. *Information and Documentation – Records Management (ISO 15489-1)*. Geneva: ISO. 2001. p. 7.

TWELVE YEARS OF USER SERVICES FOR OCEAN TOPOGRAPHY USERS: AVISO EXPERIENCE AND LESSONS LEARNED

Vinca Rosmorduc⁽¹⁾, Nicolas Picot⁽²⁾

⁽¹⁾ CLS/AVISO, 8-10 rue Hermès, 31526 Ramonville StAgne, France, vinca.rosmorduc@cls.fr

⁽²⁾ CNES, 18 av. E. Belin, 31400 Toulouse, France, nicolas.picot@cnes.fr

AVISO USER SERVICE

Aviso (Archiving, Validation and Interpretation of Satellite Oceanographic data) has been set up in 1992 to process, archive and distribute data from the NASA/CNES ocean altimetry satellite Topex/Poseidon (T/P). First goal was to distribute Geophysical data records and “quick-look” data to T/P PIs and Co-Is. Twelve years later, Aviso user service's activities now encompass:

- Operational distribution of Topex/Poseidon, Jason-1 and Envisat GDRs;
- Operational distribution of Ssalto/Duacs Near-real time products;
- Distribution of high-level altimetry products;
- A catalogue of altimetry, orbit determination and precise location products;
- Promotion of ocean altimetry, orbit determination and precise location activities.

Milestones

Major milestones were:

1992: creation of the service, launch of Topex/Poseidon

1993: distribution of Quick-look data

1994: opening of the website, T/P GDR data release

1995: distribution of level-2 products, multi-altimeter sea level anomalies and corrected sea surface height products (T/P and ERS)

1998: distribution of gridded products

2001: launch of Jason-1

2002: distribution of near-real time products

2003: opening of a Live Access Server and Opendap; Jason-1 GDR data release

All GDR data (T/P and Jason-1) were first released to PIs, for calibration and validation. Various enhancements were thus possible, with the goal of keeping the quality at its highest. Then the data were distributed to a wider audience, including reprocessing when needed to enhance the quality of the data.

Challenges in the next years will be compatibility with other datasets, and easiness of access and of use of the data. Aviso user service is building on its experience to try and face them, working hard to foster closer and closer collaboration between project teams and data users.

Organization

Aviso user service is open 52 weeks a year, on working days. Mails are centralized, with a unique e-mail address, and answered within 5 working days. Expert advices are requested whenever useful, with the whole panel of expertise available to us, from instrument to oceanography, geodesy, including Calval and processing expertise.

The user service is also closely linked to the outreach team, allowing strong interactions. Outreach targets are also potential users – among others; it is a mean of explaining the data, what they allow, and to attract attention on the technique. Aviso user service often answers high-school student questions, a service which could be presented as outreach, on the whole. On the other side, user feedback gives information that may become outreach subjects – like new research results, new applications, etc – and also help to enhance outreach documentation available.

All information, whether on data or outreach are updated on Aviso website <http://www.aviso.oceanobs.com>, in both French and English (a Spanish section is also available, for basic information)

Aviso Users

Aviso users evolved along those twelve years, in number (see Fig. 1, for GDR users, the first distributed product, but not today the most popular), but also in profile. First users were PIs/Co-Is, i.e. mostly altimetry experts or at least physical oceanographer with a working knowledge of altimetry. Now we reach a much wider audience, including

biologists, chemical oceanographers, climatologists, meteorologists, end-users, students and even schools! The need to diversify the available datasets went with the broader audience –and vice-versa.





Fig. 1. Evolution in Aviso GDR users over the years, from 35 in 1992 to over 200 today. The most popular data, MSLA, which are distributed since 1998, have about 430 registered users.

Number of countries also increased, with more and more developing countries joining the altimetry users (Fig. 2). This is an important issue, since the facilities we are used to in Europe may not be available to them (high bandwidth Internet, recent computers with DVD Rom drivers, computer engineers to help and/or program, etc.), and since those new laboratories getting involved have not always the background in data processing that most developed countries laboratories benefit from. The helpdesk function of the user service is thus much more important with those new users than even for PhD students elsewhere.



	Satellite					
	Jason-1	Topex/Poseidon	Envisat	ERS-1&2	GFO	Merged
<i>Gridded Sea Level Anomalies (MSLA)</i> (1/3°x1/3° on a Mercator grid) Study of ocean variability	DT NRT	DT NRT	NRT		NRT	DT NRT
<i>Gridded Absolute Dynamic Topography (MADT)</i> (1/3°x1/3° on a Mercator grid) Study of the general circulation						DT NRT
<i>Along-track Sea Level Anomalies (SLA)</i> Study of ocean variability	DT NRT	DT NRT	DT NRT	DT	NRT	
<i>Along-track Absolute Dynamic Topography (ADT)</i> Study of the general circulation	NRT	NRT	NRT		NRT	
<i>Along-track Wind/Wave data</i> Marine meteorology, ocean-atmosphere gas transfer studies	DT NRT	DT NRT	DT NRT			
<i>Gridded Wind/Wave data</i> (on a 1°x1° grid) Marine meteorology, ocean-atmosphere gas transfer studies	NRT	NRT	NRT		NRT	
<i>Along-track Geophysical Data (GDR)</i> Geophysical studies, operational oceanography	DT NRT	DT NRT	DT NRT			
<i>Along-track Geophysical and sensor data</i> Expert use; coastal, ice studies or anything requiring a different retracking function than the one used for ocean	DT	DT	DT			
<i>Gridded geostrophic currents</i> (1/3°x1/3° on a Mercator grid) Study of ocean variability and circulation						NRT

Table 1. Aviso products, per satellite
DT: Delayed-Time; NRT: Near-Real Time

TWELVE YEARS OF EVOLUTION

Evolution in Data Distribution Medias

When Aviso began distributing its data in 1993, CD Rom was a new thing in the laboratories, so much so that the Space agencies involved in Topex/Poseidon (NASA and CNES) decided to pay a CD Rom driver to their PIs. You then had to explain how to mount the CD driver when using VMS or Unix... CD Roms soon became a standard, to be now substituted by DVD Roms, with a higher storage capacity. But the major change in data distribution during those twelve years has come from Internet. Online distribution became possible around 1995, and feasible (at least for developed countries) about 1999-2000 for important amount of data. We want to emphasize, however, that if downloading over 100 MB is commonplace in North America, Japan or Europe today, developing countries still have problems with files over 5 or 10 MB. If we are to reach scientist in those countries, we can't rely only on Internet for data distribution, or at least we have to consider the possibility of extracting limited geographic areas.

CD and DVD Roms

Aviso delayed-time data are now mostly distributed on DVD Rom (DVD-R) rather than CD Rom. Jason-1 + T/P GDR are thus sent on DVD every 2 months (6 cycles of both missions) to registered users. These medias are quite commonplace, and do not pose any problem, even DVD Roms (however, we chose not to use double-density ones, in order to keep closer to a standard).

FTP

FTP is the oldest online data distribution protocol. It has the advantage of being easily automated, included in a processing segment, and to allow retrieval of long series of data in one shot (if the bandwidth is high enough). Such tools have made possible the distribution in real time, a sine qua non condition for operational applications.

Web-based Data Distribution

Web-based distribution is the newest “media” available. By web-based distribution,, we do not mean only information and a link toward an FTP site, but real use of the web to transport data. Interactive tools have now reached a maturity that able a center like Aviso to use them for its data. We thus opened an Opendap and a Live Access Server last year [1], [2], with some success, since about 500 domains connected to it in a year (the Opendap enables access to remote data sets through familiar data analysis and visualization packages; the Live Access Server is an interactive on-line visualization tool for plotting gridded data).

Changes in Data Format

Data format convention changed during those years. New formats appeared, more largely disseminated, and/or compacter, self-explaining or not. Now for all but the Geophysical and along-track wind/wave data, Aviso data are in NetCDF, Unidata’s network Common Data Form[3]. This matches a clear preference – even if unvoiced as such – of the users for data that they don’t have to read a handbook to use. It must be noted that ascii format have quite a lot of supporters, but was put aside for file weight questions. However, with interfaces like the Opendap, ascii can be retrieved from data that are in NetCDF.

Altimetry data are along-track data to begin with. However, this is not the easiest thing to manipulate, when you wish to study a given area, for example, and induces reading programs a bit more complicated than a grid does. Gridded data thus are the preferred “format”, by far. It does suppose, however, that some of the grid points are interpolated not measured, ones.

Changes in the Means of Communication

Changes also involved communication means. Internet has also played a major role here, both for data center-to-user communication and for user questions to the data centers. Indeed, whereas in the past you had to know about a data distribution center before you contacted it, now web sites able would-be users to search for themselves where the data they need exists, to have complete, immediate information about them, and to compare between format and services, etc.

And the change is also the other way round: users can also interact much more easily via e-mail than they did by ‘snail-mail’, when you had to take a sheet of paper (maybe even a typewriter) to ask for something. As User Service goes, it must be said that it is sometimes a mixed blessing, since quite a lot of people find it easier to send an e-mail than to have a look at the data handbook!... Still, user feedback and service enhancement are thus much more easier.

Satellites and Merging of Data

There are now no less than four available altimetry satellites: Topex/Poseidon, GFO, Jason-1 and Envisat. The two ERS satellites were used successively, and are now either shut down (ERS-1), or only usable for altimetry over a limited area (ERS-2, with no onboard recorder for the altimeter since June 2003); ERS-2 is anyway on the same track than Envisat, close to it, and so would not bring more information.

One of the main issues with altimetry today is that its data are “only” along-track, and that leads to the need of a compromise between spatial and temporal resolution when defining a mission. So, if you wish both high spatial resolution and high temporal resolution, one of the solutions is to combine data from several satellites. But, to do so, algorithms had to be developed in order to homogenize all satellite datasets, calibrating them on the highest-quality one (Topex/Poseidon at first, now Jason-1). Aviso started early in the merging of data, since we did it for the two Topex/Poseidon onboard altimeters (Topex, and Poseidon-1, which was on 10% of the time as an experimental passenger), and then began in 1995 to merged Topex/Poseidon and ERS-1 data, cross-calibrating the less accurate ERS on T/P at crossover points.

Evolution in Users

In twelve years, users themselves have also evolved. From altimetry specialists to neophytes, it has already been mentioned above in this article. From programmers to software users, in a large part, too. But another, deep, change is that oceanographers are not anymore specialized in a specific observation technique. They are now using every available observation system and also models outputs to observe and study a class of phenomena. Thus they do not wish to delve too deeply in each technique characteristics, and expect to benefit from “standardized” dataset format, and

highly processed data representing not so much an actual measurement, than a physical parameter representative of the ocean state and dynamic.

ADAPTING DATA TO USERS OR USERS TO DATA?

Adapting Data to Users

Taking User Feedback (even unvoiced one) into Account

User requests often unearth questions, problems, and/or requests for specific products. For example, we decided some time ago to stop distributing the Corrected Sea Surface Height data (CORSSH) product, on the assumption that, between GDR and SLA, the users could find what they needed. However, repeated requests for this particular product revealed the need for an intermediate product, simpler than the GDR (containing all the measurements, corrections, flags, biases, etc.), but with more data fields than the Sea Level Anomaly product (SLA, “pre-computed” data, with only longitude, latitude, time, and SLA values), in particular data where you can access total currents – i.e. where mean dynamic topography were not subtracted. This has led to two different products: Absolute Dynamic Topography, both gridded and along-track, representing total ocean circulation, and a new CORSSH product, to be released soon. Similarly, repeated requests for gridded wind speed modulus and significant wave height data have prompted us to develop such products.

Two surveys of Aviso users were conducted, one in 1995, the second in 1999. First one led to along-track T/P and ERS homogeneous CORSSH and SLA. The second was meant to define Jason-1 products. Good answer rates allowed us to improve our service, and to adapt products to the user preferences.

Discussing with users during meetings is also a valuable source of feedback that can bring forth things that people would not spontaneously express. It is also a way of feeling the mood, to prepare for future data.

We also began recently to have a closer look at articles published in science journals, with the following questions in mind: who’s using Aviso data (who’s saying they do? What is the proportion compared to other sources of similar data)? For what sort of study/applications are they used?

An Equation with Five parameters: What the Data Represent, their Format, Delivery Delay, Delivery Channel & What the User Wants to Use them for?

What the data represent is in itself an important consideration with respect to users. Altimetry waveforms will never appeal to non-experts, whereas the more processed gridded data will seem insufficient to experts, with interpolation errors, no access to each correction values... A range of products must take that into account to be complete – or close itself to some users.

Data formats are also an issue. They must be at the same time as easy as possible to use, complete, and compact. Thus a compromise has to be done. E.g., Ascii would be the easiest to use, but is quite heavy, not standardized, and not self explaining.

Delivery delay is what constrains the possibility of operational applications: data have to be available in real or at least near-real time in order to get used in nowcasts or forecasts. Such deliveries began experimentally with Topex/Poseidon, and are now included in the Jason-1 and Envisat mission specifications. But, since for altimetry data high quality need time (to compute the best possible orbit), some applications can’t use those near-real time data. Delayed time data have thus to exist, as reference and as the only possible tool when you wish to study things like mean sea level rise, with need of very high accuracy.

Delivery channel is an issue as to whom we want to give access to the data. As mentioned above, developing countries have not yet internet bandwidth capabilities to give them access to heavy online files. Decision to stop or continues DVD delivery is a somewhat political one – do we wish to spread the benefit of the data abroad, or not?

The fifth element to take into account is “what for”: all datasets can’t be used for all purposes. If Waveforms or even GDRs contain, by definition, all the possibilities of the other, higher-level, datasets, they may not be practical to users. Some uses imply non-standard processing (ice, land applications, hydrology in a certain measure): thus we need to decide whether send the Waveforms for those applications, or to deliver specific data for those uses. For example, here is a list of possible applications that have different needs:

- Geodesy
- Ocean circulation and variability

- Ocean forecasting
- Tides
- Weather forecast and studies
- Air-sea interactions
- Climate forecasting
- Global climate change
- Hydrology
- Land applications
- Ice applications
- End-users

For all those, different products could be defined, with different parameters (e.g. GDR reprocessed near land for hydrology, etc.).

Defining User Profiles

The “equation” above seems quite complex, and all combinations can’t realistically be served, especially if you try to take also into account the user capabilities (computer skills, scientific literacy...), their hardware, OS, etc. However, some combinations are not useful (or highly improbable), some are evident. For the others, compromises have to be found, by evaluating the most useful datasets corresponding to the most likely user. Short guidelines are presented with the data on the Aviso website, to try and help new users to find the data matching their profile. They should be expanded, so that any new user will find the product best fitting his/her needs.

Adapting Users to Data

Even the best user service in the world can’t use the data instead of the user! So it is important to “adapt” the users to your data. That means that you have to offer them clear, to-the-point, information on the data; not only handbooks – those are the last resource of the desperate, or for expert users... You have to guide new users with the idea that everything is really new to them, so as to have them learn to find their way through a quite complex world of data, and continue by themselves afterwards.

CONCLUSION

User services distributing datasets like Aviso have thus to choose between serving a complete range of users, increasing the complexity of their product range, or targeting specific, pre-defined users, letting the other users cope with whatever exists. Aviso will continue to propose a wide range of products, trying to enlarge our offer to best fit the needs, even if we are aware that we can’t serve the exact need of each individual user.

Future developments will be in the field of data deliveries, especially with interactive tools. Such technologies are a major issue that should be closely kept under survey by data user services, with in mind standardization and easiness of use (more than the appeal of brand-new technologies).

On another plan, joint databases or web portal with data from several available observing systems (in situ, altimetry & other space techniques) and models should be developed during the next years. Such portals could benefit users by providing them one-stop access to whatever data are available. However, even the most easy-to-use, interactively accessible and self-explaining data shouldn’t be distributed if there is nobody to answer to users’ questions. So, having people connected with data processing to answer users can’t be overlooked in such portals. Thus, user services linked to the data production center must continue their activities, in a somewhat broader data world.

In twelve years, many things have changed, in technologies, data, media, even users... Aviso user service had to adapt continuously. It was, indeed, the wish of one of its patrons, Michel Lefebvre [4] that “Aviso is not just a technical solution, important though that is, but rigor and innovation blended into constant added value”. He added “The division of labor will need to be more flexible than ever. This is the really tricky part, needing continual feedback. The hardware, the data and the partners will change, but our methods are open-ended enough to move with the times”. We hope to have met this challenge, and to continue meeting it during the next years.

References

- [1] S. Baudel, F. Blanc, T. Jolibois, V. Rosmorduc, "Discover, browse and access ocean data on internet", *this issue*.
- [2] Aviso Live Access Server: <http://las.aviso.oceanobs.com>; Aviso Opendap: <http://opendap.aviso.oceanobs.com>
- [3] Unidata NetCDF pages: <http://my.unidata.ucar.edu/content/software/netcdf/>
- [4] M. Lefebvre, “Personal view: Aviso, the first movement”, *Aviso Newsletter*, 1, p.2, July 1992.

Preserving Access to Legacy Information Through Data Migrations at NSSDC: Experiences and Lessons Learned

Donald Sawyer ⁽¹⁾, H. Kent Hills ⁽²⁾, Pat McCaslin ⁽³⁾

*⁽¹⁾ National Space Science Data Center(NSSDC)
Code 633 NASA/Goddard Space Flight Center
Greenbelt, MD 20771 USA
Donald.M.Sawyer@nasa.gov*

*⁽²⁾ QSS Group, Inc./NSSDC
Code 633 NASA/Goddard Space Flight Center
Greenbelt, MD 20771 USA
Hills@mail630.gsfc.nasa.gov*

*⁽³⁾ QSS Group, Inc./NSSDC
Code 633 NASA/Goddard Space Flight Center
Greenbelt, MD 20771 USA
McCaslin@mail630.gsfc.nasa.gov*

INTRODUCTION

During the 40 years of operations of the National Space Science Data Center (NSSDC) at NASA's Goddard Space Flight Center, NSSDC has accumulated an archive of over 4,400 distinct "data sets", of which over 2,300 are digital, from 1300+ instruments flown on 370+ mostly-NASA scientific spacecraft. The current digital data volume is about 25 TB and it resides on a number of media types, including magnetic disk, optical disk, offline tapes and tape cartridges, CD-ROMs, DVDs, and Digital Linear Tapes (DLTs) in a near-line jukebox. This wide variety primarily reflects the long history of NSSDC during which the media type on which data were sent to NSSDC also became the media type used to hold the data for the long term. Nevertheless, efforts have been made to move data from media types no longer well supported, such as 7-track tape, to then newer media types such as 9-track tapes and 3480 cartridges while still retaining key original-media information. More recent migration efforts have centered on transforming data into canonical formats and moving the resulting data packages to DLTs, while still retaining key original-media information, to improve future migrateability.

A number of issues have arisen and were tackled in ways that made sense at the time, given the level of resources available. These issues included:

- 1) the handling of computer words that don't fit evenly into 8-bit octets;
- 2) the handling of character data expressed as 6-bit or 8-bit characters;
- 3) original documentation as paper documents describing the layout on a particular type of tape;
- 4) completeness of checking original submissions for conformance to the documented formats;
- 5) formats that depend on a particular operating system's file structure, such as VAX/VMS;
- 6) consistency, or lack thereof, in populating syntactic metadata about a given data set and/or the media on which it resides.

These issues, the way they were handled, and the rapid evolution of technology have resulted in an environment that is currently sub-optimal for the most cost-effective long-term preservation of the information content. It is argued that a clear specification of the information content, in terms of data and documentation (metadata), and independence from the underlying media, is the current best approach to ensuring cost-effective preservation of this science information. It is also necessary to make the information practically useful to current researchers. Bringing all of NSSDC's legacy data to this state is hampered by a past lack of consistency in the processes employed, and in the documentation and validation of these processes. These problems are being overcome by making the effort to ensure that current migration efforts result in data packages having the desired characteristics (including accurate syntactic/format information) and by ensuring the documentation of archival decisions and processes.

HISTORICAL PERSPECTIVES

Decade of the 60's

The National Space Science Data Center (NSSDC) was the first official archive established by NASA. Formed in 1964, it was made responsible for acquiring and preserving the data and documentation obtained from, and/or derived from, space flight instruments. While the majority of the data sets are from NASA missions, NSSDC also archives data from other US agencies and from foreign agencies. Significant amounts of digital data began to be received in 1967 and were stored on 7-track, 200 or 556 bits-per-inch, magnetic tape. Some analog recorded data on magnetic tape were also received. Approximately 220 digital data sets were received on over 2000 magnetic tapes during this decade. Computer programs were typically held on punched paper tape or punched cards. There were also significant amounts of data on various types of film and paper. The film archive will not be discussed further in this paper. Virtually all documentation was provided on paper.

Data from satellite instruments were generated in very compact formats to minimize the bandwidth needed to transmit them to the ground. Data fields composed of individual bits, or groups of bits, were used to represent various data values, and these were typically multiplexed in a varying data structure that depended on the frequency with which the data fields were sampled by the science instruments. On the ground, a variety of computer systems was used to process the complex data structures, resulting in a large variety of computer word, data field, character, and floating point representations. There was no manual on how to manage an archive of such data, so NSSDC proceeded as seemed best. Magnetic tapes were duplicated, and software was written to record and manage the information about the data. NSSDC also served the science community by publishing a record of its data holdings and by providing copies of tapes, along with hardcopy documentation, to data requestors. Extensive procedures were documented for the processes of acquiring, managing and serving the data. This was supported by the use of paper forms, filled out by Acquisition Agents (scientist with appropriate expertise), that were given to key-punch operators. They produced punched cards that could be read and incorporated into the information system. A key-punch verifier was used and consisted of a second key-punch operator who re-keyed the data that was simultaneously compared to the key-punched cards. An IBM 7094, installed in 1967/68, was the heart of the main processing system. It read and wrote 7-track tapes, and it input programs and supporting data from punched cards.

Decade of the 70's

The situation during the 1970s was not substantially different. At the beginning of the decade, the major components of the NSSDC hardware environment included the IBM 7094, six 7-track tape drives, a disk drive, a card punch, and a card verifier. There were over 2000 1/2 inch magnetic tapes that grew to over 60,000 tapes holding an additional 630 digital data sets. While some 9-track tapes were received, these were processed on another machine at the Goddard Space Flight Center. In 1975, a MODCOMP IV computer replaced the IBM 7094. This provided significantly greater computing power and memory. It processed instructions at about one million per second, had 0.5 megabytes of memory, and four 50-megabyte disk drives.

To obtain a sense of typical operations of the time, the processing of incoming data sets was nominally handled in the following manner (quoted and paraphrased [1]).

"All magnetic tapes received by NSSDC were first entered into the storage records by filling out the proper paper forms and assigning a unique accession number. At this point an acquisition agent, to be assisted by a programmer as necessary, was assigned to the data set for preliminary analysis. The following steps were taken as appropriate.

1. Data and documentation reviewed
2. Processing needed was determined
3. Entries made into the information system
4. Receipt of data set acknowledged (letter) and discrepancies noted.

The objectives in the preliminary processing were: (1) ability to read the entire tape in its logical record format; (2) ability to list out any function or special record; (3) ability to detect errors (logical and physical); and (4) verification of the acquisition agent's understanding of the contents.

During this preliminary processing, the programmer wrote all the necessary routines to manipulate the data and reformat it, if necessary. These programs were entered into the Computer Program File which consisted of listings of complete programs and an index sorted alphabetically by content. The preliminary analysis stage was

completed when NSSDC had the ability to use and interpret all data in the data set. This often required additional contacts with the science investigator who submitted the data.

At this point, the acquisition agent and programmer defined the format of a document called the Data Set Catalog, the functions of which were to:

1. Provide an index to the contents of the data set
2. Provide a series of error checks to be made on all tapes of the data set
3. Calculate bounds or distributions of functions using the data values found
4. Provide a useful tool to the request agent for identifying data
5. Provide a coarse description of the information in the data set

Some items in the catalog were: description of the record format, ranges of important data fields, tape-oriented descriptions of major functions, tables of gaps and overlaps, etc.

After the requirements for the catalog were defined, the programmer wrote a program to produce the Data Set Catalog. The program should also produce a copy of the original tape or a reformatted version. After checking the program and turning it over to the computer people, the rest of the tapes in the data set were processed.

The Data Set Catalog was considered available for distribution only after the following steps had been taken.

1. Table of contents is prepared, material is bound (paper), and catalog is stored
2. Program is entered into the Computer Program File
3. Information on all accession numbers has been entered into the tape tracking information system
4. All other information required has been entered into the information system"

This process of handling incoming data was quite labor intensive. There were virtually no data standards to rely on and the information systems used were developed in-house. At the beginning of this decade, there were several separate information systems employed, including:

1. Automated Internal Management (AIM) system to hold primary information about spacecraft, experiments, and the data sets that had been received.
2. Technical Reference File (TRF) to assign unique identifiers to received documents relating to satellites, experiments, and data sets. and to maintain bibliographic information for searching.
3. Request Accounting Status and History (RASH) to hold information about requests and their processing

These information systems were supported by several other types of files including the Magnetic Tape Unit Records (MTUR). MTUR assigned an accession number to each magnetic tape and associated a storage location in the archive. The overall processes were well engineered and complete for the technology available at the time.

Some of the major challenges during the 1970s were:

1. Obtaining adequate documentation and reviewing this for understandability and completeness. The documentation obtained was hardcopy and sometimes described multiple data sets from the same experiment.
2. Writing software to decode the data enough to check it against the documentation. Often this was done only for a few fields in a given file on a tape. Often it was done just from an octal dump of the records. The variety of computer words and corresponding floating point and integer representations result in complex documentation.
3. Managing the many boxes of tape that frequently showed up, unannounced, at the loading dock. Sometimes there were only the initials of the shipping clerk as a clue to where the data came from.
4. Updating the information systems with information and satisfying new requirements, as they evolved.
5. Inspecting and testing the magnetic tapes used to hold the data as it became known that there was a wide variety in the quality and permanence of tapes, even from a given manufacturer.

In an effort to obtain independence from the media on which the data arrived, and to make it easier to process data for requesters, an effort called 'Standard Tape' was initiated. The idea was to reformat incoming tapes into a standard form. However this was ultimately abandoned as the variety of incoming data structures and relationships exceeded the standard tape data model capabilities.

Decade of the 80s

NSSDC started the decade with over 60,000 1/2 inch magnetic tapes that grew to over 113,00 tapes, twenty-eight floppy disks, and 51 optical disks holding 630 additional digital data sets. The MODCOMP computer now had two 7-track tape drives, four 9-track drives, and eight 50 MB disk drives. Modems were used to provide access for running software and for providing remote (e.g., across the Ocean) displays for an activity called Coordinated Data Analysis Workshops. These were group activities focused on analyzing data to address scientific problems, and they were effective in bringing in many small new event-oriented data sets to support each workshop. In 1984, NSSDC acquired a VAX 11/780 computer. At about the same time, a new MODCOMP (Classic II) was obtained to replace the aging MODCOMP IV for magnetic tape processing. Other computational functions were transferred to the VAX, and NSSDC began to use the VAX network capabilities to entertain electronic requests for data to make some data, and supporting metadata, network accessible. A VAX 8650 replaced the existing VAX in 1986.

During this period, the metadata about data and about NSSDC operations, including requests, were kept in information system files. These were managed as a set of about 13 separate, hierarchically organized, file systems under an application called SIRS (System for Information Retrieval and Storage). There were now seven major file systems as follows:

AIM: Automated Internal Management File- 3-level file containing information on all spacecraft launchings, onboard experiments, and resulting data sets which are recorded or archived at NSSDC.

MISC: Miscellaneous Microfilm File - 2-level internal inventory file of microfilm data sets at NSSDC.

NSDF: NSSDC Supplementary Data File - 3-level file containing information on space science supportive data which are available from NSSDC and which are not data from a single satellite experiment.

RAND: Request Activity and Name Directory File - 2-level file containing information pertaining to the names and addresses of NSSDC data users and information about the data requested.

TAPE: Tape Inventory File - 1-level internal inventory file of digital magnetic tapes at NSSDC.

TRF: Technical Reference File - 2-level file containing information on documents of interest to NSSDC and its users

VTAB: Validation Table File - 2-level internal file containing acceptable values for given fields in other SIRS files. Used as a validity check during the SIRS update process, and as a table look-up during the SIRS print process.

In 1987, the first major migration of data from older magnetic tapes was begun and was referred to as the Data Restoration Program. This initial migration involved rewriting 4.5 GB of IMP (Interplanetary Monitoring Platform) 8 magnetometer experiment data from 115 tapes onto 31 full, 6250 bpi, magnetic tapes. NSSDC also purchased 12-inch optical disk drives and platters to provide greatly increased, randomly accessible, storage for data. NSSDC began tests to define and resolve various technical issues in the copying of large numbers of (possibly) deteriorating magnetic tapes. As reported in [2], "Tapes from the various data sets under analysis are being sampled to establish typical read error rates and procedures to optimize the data recovery rate with respect to the throughput of 'restored' tapes. Ongoing work includes the analysis of given tape formats as to their sensitivity to the loss of specific tape records, improving procedures to best recover important tapes that have become highly deteriorated, and testing the use of alternative Goddard facilities outside of NSSDC itself to improve throughput." It went on to say, "Data are currently being copied to high-density standard (6250-bpi, 9-track) tapes. NSSDC is exploring an expanded use of optical disk and various cartridge tape/helical scan technologies as well. The gradual transition to a full-scale production mode for restoration activities is expected over the next several months." By the end of 1989, data from approximately 10,000 tapes had been migrated to about 1,000 new tapes with copies on 3480 cartridges.

Decade of the 90s

NSSDC entered this decade with over 113,000 1/2 inch magnetic tapes, twenty-eight 5 1/4" floppy disks, and 51 optical disks that changed into about 60,000 9-track tapes and a variety of other digital media holding an additional 730 digital data sets. It also had thousands of 3480 cartridges as backup copies of data migrated to 9-track, 6250 bpi, tapes.

NSSDC also developed an archive and distribution service called NDADS (NSSDC's Data Archive and Distribution Service). This consisted of two VAX 8250s and one VAX 6410 in a cluster configuration with almost 1.2 terabytes of online optical disk storage housed in two jukeboxes. The information about the data held in this system was maintained in a SYBASE database data management system. Subsequently a new VAX 9410 was added to NSSDC's computational capability, augmenting the previously existing VAX 8650. The GSFC IBM 3081 was also used to support some of NSSDC's tape copying and migration efforts, including the generation of 3480 backup cartridges.

By the end of 1992, NSSDC's principal computing resources were a VAX 9410, three VAX 6410s, a VAX 8650, MicroVAXes, a MODCOMP Classic computer, a Britton-Lee IDM 700 data base server, and a few DEC, SUN Microsystems, and SGI UNIX workstations.

The status of the data migration and restoration effort, begun in earnest in 1988, was described in [3]. At the start, NSSDC held about 80,000 unique tapes grouped into about 1,500 distinct data sets. These ranged in size from one tape to a few thousand tapes. The 9-track, 6250 bpi tapes and 3480 cartridges accommodated about 160 to 200 Mbytes/volume. By 1993, NSSDC had migrated 24,600 tapes onto 3,900 new pairs of data volumes with a recovery rate in excess of 98%. (Note: NSSDC does not consider this an acceptable loss rate.) It was estimated that this restoration effort was costing about \$20 per input tape, with 88% being labor costs and the rest media costs. Computer time was not included, and no re-formatting was done. However about 1/3 of the tapes were 7-track, and when writing to 9-track tape, the extra two bits per byte were padded with zero's. The overall conversion of 7-track data was slow, averaging about 100 tapes per week. To speed this process, NSSDC developed an 80386 MS/DOS based system to read and write 7 and 9-track tapes. However some tapes, from the mid-1970s and from specific vendors, were classified as 'sticky tapes' in that they tended to cause buildup on the tape read heads. This was less of a problem on the MODCOMP due to a lower tape speed, so these tapes were handled by this machine. After a copy of an original tape was made on a 9-track tape, multiple original tape images were stacked onto a single 9-track, 6250 bpi tape. The 3480 copy was made of this stacked tape. Using the new equipment and an adjustment in procedures, the average migration rate went up to 200 tapes per week. There was a verification of the new tape's content and quality, and then usually the old tapes were re-cycled. Entries were made into the NSSDC information file regarding the new tapes, including which original tapes were used. Finally, one of the copies of the new tapes was sent off-site for safe storage.

Also by 1992, the major NSSDC information system files had been migrated to a relational system called RSIRS (Relational SIRS). These were the tape inventory (now called IDA), IRAND (Interactive RAND), AIM, TRF, and NSDF (now just NSD). This involved incorporating the validation table information (VTAB) into the relational scheme as appropriate.

By the Spring of 1995, NSSDC had migrated 35,000 tapes to 6,100 pairs of 9-track and 3480 cartridges. This concluded the data restoration program as no tapes were then more than 10 years old. The plan was to do continuous data preservation from this point on to ensure the accessibility to the data. NSSDC purchased a Digital Linear Tape (DLT) jukebox with a 2.5 TB capacity to augment its mass storage capacity.

During this decade NSSDC also began to transfer its Earth Science data to specific Earth Science archives. Over 2,000 9-track tapes and over 1,900 3480-cartridges were transferred to various archives, while about 26,000 9-track tapes still waited transfer at the end of the decade. It was determined that NSSDC would hold these later tapes until the GSFC Earth Science group was ready to take them.

Current Decade

NSSDC started the current decade (2000) with a variety of digital media types including over 60,000 9-track tapes (about 26,000 belonging to Earth Science), 9, 868 CD-Write Once, 1,148 CD-ROMs, over seven hundred 8 mm tape cartridges, over seventy 4 mm tape cartridges, over five hundred 12" WORM (Write Once Read Many) platters, 168 magneto-optical disks, and 163 DLT cartridges.

The number of 9-track tapes was down from 1990's initial count due to the data restoration activity that included the use of higher density tape. The growing diversity of media for long-term storage increases management costs and is being addressed by additional migrations. The data on the WORM platters were mostly accessed by the NDADS optical

jukebox system. An additional 145 digital data sets have been added this decade, mostly on CD-Write Once and DLT cartridges.

As maintenance support for the optical drives was disappearing, it was necessary to start a new migration. A system called DIONAS was developed. Advantage was taken of the Reference Model for an Open Archival Information System (OAIS) [4] by re-engineering the NSSDC ingest and archival storage system. Key changes were the definition of an Archival Information Package (AIP) that includes metadata along with the data files in a canonical form independent of the underlying media, and the use of two DLT jukeboxes with an associated Oracle database to provide long-term storage and preservation. This system, and the start of this migration, were described by several articles [5]. A major driver for both the system design and the migration process was schedule, as the hardware supporting NDADS was soon to be decommissioned. The system design and the processes used during the data migration reflect the perceived urgency of a quick data transfer.

The migration of data from the VAX/VMS-based NDADS system to a UNIX-based system involved solving the problem of retaining file and record attributes maintained by the VMS system but not by a UNIX system. This was done by capturing key attributes and packaging them with the data in AIPs, and by reversibly transforming some data sets to incorporate record delimiters within the data stream. During the process of testing this approach, it was discovered that some initial transfer errors had taken place when some data sets were originally entered into the NDADS system. These were traced primarily to inappropriate usage of FTP (File Transfer Protocol) in text mode that somehow evaded the subsequent quality checking. Fortunately these errors could be corrected by appropriate algorithms that were executed as the data were migrated from NDADS to DIONAS. However another type of error was discovered well into the migration. For a particular file type, the VMS system generated a spurious end-of-file that caused the data to be truncated. Data validation during the migration was not sufficiently rigorous to catch this early in the process and these data had to be re-migrated over an extended period of time.

Another problem occurred during the migration when some optical platters were written directly to DLT cartridges so as to allow the timely retirement of the NDADS jukebox. Backups of these cartridges were not made. In this mode, the completion of the migration was from these DLT cartridges into DIONAS. One of these DLT cartridges was subsequently overwritten due to operator error and extensive efforts were required to re-obtain the original data from the original suppliers. The migration of the data from NDADS to DIONAS was essentially completed by December 2002. In hindsight, NSSDC might have been better served by first preserving the NDADS-resident data and documentation on available hardware using well-documented processes. The canonical data transformation into AIPs could have been effected as a second step, performed with well-planned and documented processes.

More recently NSSDC has developed plans and systems to begin the migration of its legacy 9-track tape data into the DIONAS system. This has involved a full description of the process in all its details. Much of this data was involved in the migration begun in 1987 and completed in 1995, as described previously. The approach taken has been to map the contents of an original tape, as it existed prior to the 1987 data restoration effort, into an NSSDC Archival Information Package. This has required a look at the existing documentation for each tape, or data set, to ensure that the tape and format description matches reasonably with the format found today. In particular, this is primarily a look at tape block sizes, at least initially. However because some of the original data came from 7-track tapes and had been padded at the byte level (6-bits into 8-bits), the new block sizes will not match the original documentation. NSSDC also found that not all documentation was updated to reflect this change. A case-by-case analysis is required when discrepancies are found. Further, to make this migration maximally efficient, it was planned to use the IDA (tape inventory) database in an automated pipeline approach. Unfortunately, enough errors and omissions have been found in this data base that it is necessary to transfer and update the needed information into a second database that will drive the migration. As of July 2004, this migration is about to go into production mode, so additional insight may be expected in the near future.

Figure 1 shows a high-level schematic of the data sources, flows, processes, and stores. Check sums are generated for each file by the Media Read process and are used to ensure correct final results. The Data Migration Database manages the first stage of migration leading to the generation of AIPs. The second stage of migration is managed by DIONAS and results in AIPs on DLTs along with copies of data files and attributes on an Anonymous FTP site.

The design of this latest tape migration has also brought into focus a need for a more consistent view and handling of the metadata or documentation about the data, and about the processes of managing the data. The various information systems, delineated previously, are not yet fully integrated. This reflects the need to review and re-engineer the overall data flow and metadata capture, and update processes, at NSSDC. This is an ongoing effort. The IDA tape inventory database will be replaced by JIN (Java-based Inventory) and the AIM database has been incorporated into the JEDS




Fig. 1. Legacy tape migration data flows and processes

(Java-based Experiment, Data, and Satellite) view of the overall NSSDC Information Management System (NIMS). Efforts are underway to provide digital copies of all the significant legacy data set documentation and to make them the official source. Additional updates to the information system are needed to capture both the design of this latest migration effort and the experience as it unfolds.

EXPERIENCE SUMMARY

Over the 40 years of NSSDC's history, the technology changes have been profound. About the only consistent digital artifact has been the bit. Since technology change has accelerated in the recent decade, there is no telling what the next 40 years will bring. Against this backdrop, it seems clear that NSSDC management has consistently promoted reasonable top-level approaches to acquiring and preserving the digital information. Nevertheless, there have clearly been problems at the more detailed implementation level that have become apparent from recent data migration efforts.

The extensive procedures of the 1970's were rigorous, but also highly manually intensive as dictated by the technology of the time. During the 1980s, the variety of computers being used and their expanding role in managing digital information led to substantially increased data and data management complexity. It appears, in hindsight, that the initial rigor was unable to be maintained. This is evidenced by the finding of inconsistencies and omissions in the IDA (tape inventory) database as it reflects the results of the 1987-1995 data restoration effort. Further, much of the data set documentation remained as hardcopy, and key information, such as blocksize, was not captured in IDA to facilitate more automated checking. It also appears that a detailed and adequate process for this first tape restoration/migration was not formally defined and followed rigorously. The same was true for the loading of the NDADS jukebox. Rather, significant reliance was on skilled personnel. This would have been satisfactory if more complete testing and validation of the results could have been performed. This trend continued into early 2000 as regards the migration effort from NDADS to DIONAS. There was extensive pre-testing and automation involved in this effort that uncovered numerous

issues and corrections. However unexpected sources of error, such as the way VMS handled a certain file type, did not show up in the pre-testing and it resulted in truncated data that could only have been found by greatly increased checking of before and after results. Note that overall checksums for data were not employed prior to the creation of AIPs in 2002, and then only during the creation of the AIPs. Had they been available on the original data, the truncation problem would have been caught immediately.

The migration of data into DIONAS from NDADS also showed the possible consequences of not making backup copies, even when experience had shown that the originals were highly reliable. This reliability was compromised by human operator error, which is always a potential problem.

The legacy tape migration design and implementation has also been a valuable experience. This has enabled a fresh perspective on the results of the 1987-1995 tape migration and the maintenance of related metadata. The lack of adequately consistent metadata to drive this legacy tape migration from existing databases is driving the generation of detailed design and process specifications in an effort to ensure that all significant metadata will be captured digitally and that all processes will be understandable not only to those managing this migration, but to future managers and operators as well. Great attention to detail is needed to achieve these objectives.

LESSONS LEARNED

1. The importance of thorough end-to-end process planning for carrying out migrations or other significant data management activities can not be overemphasized. Processes must be planned, documented, and reviewed at multiple levels in the organization. Taking shortcuts may look cost-effective but will be expensive in the long run.
2. If a problem can occur during migration or other data handling, it probably will at some point. Therefore rigorous testing of before and after is needed to identify when such problems do occur. It is not possible to identify them all in the planning or design stage. For example, the use of checksums applied to the data files as soon as possible in the data life-cycle looks prudent.
3. If an original set of data on a medium is not duplicated immediately and both verified to be correct, with one going to a 'safe' location, then at some point the single medium vulnerability will result in lost or corrupted data. Therefore the data handling process needs to address this explicitly.
4. The importance of maintaining both a current architecture of data flows and processes, and a corresponding data model of the metadata about the archive content, is critical to long-term cost-effective data management. The architecture sets the larger picture, and the data model typically gets translated into data base schema. Keeping these up to date, in the face of changing requirements and resources, can require a significant expenditure of resource. For example, if some fields are no longer to be populated, or populated only occasionally, this must be reflected in the data model and then in the implementations. Adhere strictly to database field definitions -- do not let short-term pressures overturn this rule. If a new field is needed, insert it now, rather than later, despite objections from support staff. Such changes need to be documented and readily understandable if the resulting data bases are to be relied on to drive future migrations or other major, automated, processing activities.
5. Use automation wherever and whenever possible; it reduces typographical errors and other human mistakes. It also reduces the workload on a staff that is too small already. Continually or frequently review all manual operations to see what can be automated. There may be new database capabilities (triggers, etc.) that become available so that additional steps can be automated. The scientific staff must also maintain a certain level of awareness of technology, to see what can be used to good advantage.
6. Avoid generating multiple databases that contain duplicative information, or utilize an interface that populates all databases, making the multiple databases seem as one to those populating it.

REFERENCES

- [1] N. Karlow and J. I. Vette, *Flow and Use of Information at The National Space Science Data Center*, NASA/NSSDC 69-02, pp. 19-21, January 1969.
- [2] C. Ng and R. E. McGuire, "Pilot Program Initiated for Restoration of Data on Old Magnetic Tape," *NSSDC News*, Vol. 4, No. 1, Spring 1988, p. 13, 1988.
- [3] J. King and R. Post, "More Data Restored to New Media," *NSSDC News*, Vol. 8, No. 4, Winter, p. 1, 1992.
- [4] *Reference Model for an Open Archival Information System (OAIS)*, Recommendation CCSDS 650.0-B-1, Issue 1, Blue Book, Consultative Committee for Space Data Systems, <<http://www.ccsds.org/CCSDS/documents/650x0b1.pdf>> (alternatively ISO 14721), January 2002.
- [5] *NSSDC News*, Volume 16, Number 4, <http://nssdc.gsfc.nasa.gov/nssdc_news/dec00/dec00_toc.html> December 2000

The Lifecycle of NASA's Earth Science Enterprise Data Resources

Kenneth R. McDonald⁽¹⁾, Richard A. McKinney⁽²⁾, Timothy B. Smith⁽²⁾, Robert Rank⁽³⁾

*⁽¹⁾NASA/Goddard Space Flight Center
Mail Code 423, Greenbelt MD 20771 USA
(ken.mcdonald@nasa.gov)*

*⁽²⁾USGS/EROS Data Center/SAIC
Sioux Falls, SD 57198 USA
(mckinney@usgs.gov), (tbsmith@usgs.gov)*

*⁽³⁾NOAA/NESDIS
5627 Allentown Road #100, Suitland MD 20746 USA
(Robert.Rank@noaa.gov)*

Introduction

A major endeavor of NASA's Earth Science Enterprise (ESE) is to acquire, process, archive and distribute data from Earth observing satellites in support of a broad set of science research and applications in the U. S. and abroad. NASA policy directives specifically call for the agency to collect, announce, disseminate and archive all scientific and technical data resulting from NASA and NASA-funded research. During the active life of the satellite missions, while the data products are being created, validated and refined, a number of NASA organizations have the responsibility for data and information system functions. Following the completion of the missions, the responsibility for the long-term stewardship of the ocean and atmospheric, and land process data products transitions to the National Oceanic and Atmospheric Administration (NOAA) and the U.S. Geological Survey (USGS), respectively.

Ensuring that long-term satellite data be preserved to support global climate change studies and other research topics and applications presents some major challenges to NASA and its partners. Over the last several years, with the launch and operation of the EOS satellites and the acquisition and production of an unprecedented volume of Earth science data, the importance of addressing these challenges has been elevated. The lifecycle of NASA's Earth science data has been the subject of several agency and interagency studies and reports and has implications and effects on agency charters, policies and budgets and on their data system's requirements, implementation plans and schedules. While much remains to be done, considerable progress has been made in understanding and addressing the data lifecycle issues.

Lifecycle of NASA's Earth Science Data

One way to view the life cycle of remotely sensed Earth science data is to identify the major functional areas or facilities where the data resides and is transformed during its lifetime. The breakdown that we will use includes mission planning, development and operations, science data production, active archive and long-term archive. While the physical configuration and organizational responsibilities of these facilities can vary from mission to mission, all of the functions are required. Following is an overview of the responsibilities of each area.

Mission Planning, Development and Operations

The mission planning and development functions have traditionally focused on the lifetime of the mission, from conception, to launch, through data reception and production, to and including insertion into the active archive. It includes the planning, design, development, funding, and any other aspects of preparation. After the pre-launch, launch and a commissioning period, the mission will transition into its operational phase. Typically, a mission's life span is only a few years after which the mission ends to allow for a final reprocessing of the data. The mission operations functions include both the command and control of the spacecraft and the instruments and the acquisition and low level processing of all of the data and its delivery to the science processing and active archive facilities. In addition to the science data streams, all of the orbit, attitude and engineering data are also delivered. Mission operations also include processing the raw data into a form usable by the science data users (Level 0).

Science Product Generation Responsibilities

The science product generation receives the Level 0 data and applies various algorithms to produce high-level products in support of land, oceans, atmosphere, hydrology, and other scientific studies as well as a wide range of applications (e.g. agriculture, disaster management, environmental health) that can use the science products. The generated products are distributed primarily to the active archive centers but also can be distributed directly to users.

Active Archive Responsibilities

The active archive is the system for processing support, archiving, documenting, and distributing the data and information for the life of the mission. The active archive's major role is to serve the day-to-day needs of the mission. At this stage in its life cycle, the data are being regularly processed and reprocessed using on-going data validation and quality assessment results while also being provided to the broader community of science and applications users. The role of the Active Archive typically involves three components: scientific stewardship, customer service, and IT infrastructure requirements. The active archive supports the routine operations of data acquisition, data processing, data re-processing, and data staging for archive products requested for real-time and historical data from the mission. The active archive keeps track of the various processing algorithms that may be involved in the project and the appropriate metadata and browse links within the system. It is the system hub for the mission's product generation system and it is the fundamental source of information and data for Long Term Archive (LTA) transfer activities.

LTA Responsibilities

The LTA is the archive in which stewardship of data, products, information, and documentation is held on a permanent basis or until the data is considered to be of no value. The stewardship entails preservation, maintenance, and access of the data to ensure integrity and quality of the data as the documentation indicates. The LTA is generally populated from the active archive.

A Call for Action

NASA Policy Directives specifically call for the agency to "collect, announce, disseminate, and archive" all scientific and technical data resulting from NASA and NASA-funded research. Various scientific and policy-making groups have reviewed and defined the requirements for essential data systems and services needed to ensure an LTA for satellite data record in support of climate research. Most significantly, the U.S. Global Change Research Program (USGCRP) evaluated the purpose of an LTA program for Earth observation data and derived products, lessons learned from current and past experience, and the vision, guiding principles, and essential functions necessary for the success of such a program. The USGCRP concluded [1] that a national LTA program must:

- Enable and facilitate the best possible science and highest-quality assessment for making policy and business decisions
- Document the Earth system variability and change on global, regional and local scales, building and maintaining a high quality base of data and information and establishing the best possible historical perspective critical to effective analysis and prediction
- Ensure archive holdings, facilities, and services are actively promoted and made readily available to the maximum number of users
- Enable and facilitate future research.

A report from the National Research Council (NRC) Committee on Earth Studies [2] has also contributed to the definition of requirements for essential data system services need ensure a long-term satellite data record in support of climate research. The CES report defined following principles to help ensure the preservation of the climate record:

- Accessible and policy-relevant environmental information must be a well-maintained part of our national scientific infrastructure.
- The federal government should 1) provide long-term data stewardship, 2) certify open, flexible standards, and 3) ensure open access to data.
- Because the analysis of long-term data sets must be supported in an environment of changing technical capability and user requirements, any data system should focus on simplicity and endurance.

- Adaptability and flexibility are essential for any information system if it is to survive in a world of rapidly changing technical capabilities and science requirements.
- Experience with actual data and actual users can be acquired by starting to build small end-to-end systems early in the process.
- Multiple sources of data and services are needed to support development of climate data records.
- Science involvement is essential at all stages of development and implementation.

More recently, the Earth Observing System Science Working Group on Data [3] offered the following recommendations relevant to the Earth Science Data Lifecycle:

- Science stewardship: NASA, in conjunction with NOAA and USGS, should determine what it is and how it works, at the various stages in the life of the LTA, who is responsible for it, and who funds it.
- The transfer of data sets to the LTA should begin as soon as feasible.
- Requirements of the individual instruments for the LTA should be communicated knowledgeably to NOAA and USGS.
- Coordinated schedules and goals for working with NOAA and USGS to effect the initial LTA agreements, planning, and transfer.
- Instrument teams and DAACs should participate in advisory panels and committees within NOAA/USGS to specify and administer the LTA.
- EOS teams should identify and recommend LTA levels of service that should be provided by NOAA and the USGS.
- Provision is required for NASA investigators to have ongoing access to the data sets under similar conditions to the present.

Data Lifecycle Studies and Analyses

In early 2002, NASA Earth Science Enterprise (ESE) initiated a series of studies under its Strategic Evolution of ESE Data Systems (SEEDS) activity. The purpose of the studies was to address particular topics that would help the enterprise evolve its data systems capabilities in the future. One of the topics identified was the long-term preservation of the NASA ESE data collections. While the motivation for this study arose in response to concerns with the long-term preservation of the science data and it is true that many of the associated challenges are related to the transfer of data from the active archives to the LTAs, it became clear that effectively satisfying the LTA requirements levies new or raises the importance of existing requirements on the other functional areas. Therefore, the study was broadened to be the SEEDS Data Lifecycle (DLC) Study.

The approach taken by the SEEDS DLC study team was first to thoroughly review all of the documentation relevant to data preservation that had been produced by previous study teams and review committees. This also included the review of mission and project requirements documents and relevant interagency Memoranda of Understanding. Then the study team generated a preliminary set of responsibilities for each of the data lifecycle functional areas associated recommendations. This draft study report formed the basis of the team's interactions with the participants of the several SEEDS Community Workshops in 2002-2003. The participants in these workshops represented data producers and users from Earth science research and applications communities, data systems and especially data management specialists and members of NASA's partner agencies and organizations. The input from the community interactions was incorporated into the DLC studies findings and recommendations (<http://eos.nasa.gov/seeds/>).

The set of recommended responsibilities that were identified for each functional area to ensure the long-term preservation of NASA Earth science data is too lengthy to repeat in this paper. In general, they provided detailed guidelines with respect to activities of planning, documentation and communication and the interactions between them. With the diverse set of organizations that are responsible for data over its lifetime and lengthy timeframes involved, it is essential that a comprehensive plan be developed and documented so that all parties are aware of their respective responsibilities. As data and its associated software move through its lifecycle, it is essential that all associated documentation be produced and conveyed with it to ensure that the data is not only preserved but also usable and well understood. Finally, all of the organizations responsible for the data must communicate among themselves and their user communities to ensure a smooth transition of the data and a set of data services responsive to users needs. These interactions must also be documented in the form of interface control documents among functional elements and user requirements and system operational concepts.

In addition to the specific responsibilities delineated for each functional area, the following general policy recommendations were also outlined:

- There should be an archive defined for each data product. Identify the active and LTA at the beginning of the mission, noting that for some products the active and LTA may be the same.
- "Data-buys" or any proprietary mission needs a time period to be defined after which the data becomes the property of NASA. Intellectual property rights issues need to be addressed.
- During the proposal process, the new project needs to commit/agree to a plan for a data lifecycle strategy for the project data.
- A process should be established for the science assessment of the long-term potential of data and data products to assist in making decisions on data stewardship in a resource-constrained environment.
- All archive data collections should be complete, including the archiving of the required ancillary data, project and data set documentation, and the science production software.
- Once a physical transfer of any data has occurred and been formally accepted by the archiving site, these data become the responsibility of the accepting party.
- Data may be transferred to other archives such as the National Archives and Record Administration (NARA) when deemed appropriate.
- Data should be available throughout its life cycle without loss or degradation in quality.
- Throughout a product's life cycle, a point-of-contact should be provided that could answer questions about the data or use of the data.
- Determine the relevance of Data Quality Act (67 FR 8452) and resolve issues related to watermark, provenance, reproducibility of data, peer review, integrity of data, and supporting information.

A final recommendation that was principally developed as a result of the community interactions was that NASA needs to form and maintain a Data Lifecycle Working Group with broad representation from data producers and users and data management specialists from across government agencies, academia and industry. It is clear that with the timescales that are being considered, even when the aforementioned data lifecycle plans are developed and base-lined, they will likely need revision as agencies' priorities and budgets as well as users' needs and expectations change. Some sort of working group to assist in the studies and analyses to support NASA and its partner agencies in addressing those changes will be essential.

Current Data Lifecycle Activities

NASA

Over the past fifteen years, NASA has been designing, developing, operating and evolving the Earth Observing System Data and Information System (EOSDIS) to acquire, process and archive its Earth science data resources. Since 1994, an operational prototype of EOSDIS (Version 0) largely based on existing components has been handling the heritage data collections of the agency. The data are archived at nine Distributed Active Archive Centers (DAACs) across the USA that were selected on the basis of their existing expertise in producing and managing Earth science data for particular science disciplines. The EOSDIS Core System (ECS) was installed at four of the DAACs in 1999 to support the Landsat 7 and Terra missions, the latter being the first of the EOS platforms. Since then, EOSDIS has been performing the mission operations, science processing and active archive functions for those and most of the subsequent missions of the EOS era (a select few missions have performed some or all of those functions outside of EOSDIS).

Today, EOSDIS is fully operational and handling an unprecedented amount of data from numerous instruments on multiple missions. The majority of that data is archived in the ECS DAACs, which with new data and reprocessing campaigns, are ingesting approximately 4.5 TB per day and have a collective near-line and on-line archive that is greater than 3 PB. The system also has a large and active user base. Last year, all nine of the DAACs satisfied over 220,000 users orders for data and distributed in excess of 440 TB. In addition to distributing the data, the DAACs also have a full contingent of trained user support personnel to assist in the acquisition and use of the data.

Over the course of its development EOSDIS architecture has evolved, but the major functions and the associated requirements have remained intact. As an example, in its original design, the EOS science data production was to be done by the ECS at the DAACs using science-processing software produced by the instrument teams. Today, most of this processing is done by Science Investigator-led Processing Systems (SIPS), usually co-located with the instrument teams, and the products are delivered to the DAACs for archive and distribution. However, in addition to delivering the

science data, the teams are also responsible for delivering software and documentation that are essential in assuring the long-term preservation of the data.

Understandably, the focus of NASA to data has been on getting a stable, operational system in place to support the mission operations, science data production and active archive functions for the EOS missions and that has been accomplished. However, the requirement for transitioning data to its long-term archive is the life of the mission plus four years, which in the case of Terra would be about 2010. To meet that schedule, a general guideline has been to begin to transfer data three years after its acquisition. For a number of reasons this guideline has not been met, but the issue has risen in priority and NASA is actively working with its partners to accelerate the process.

An overview of the process that is being defined to govern the transfer of data from the active archive to the LTA and which also identifies the issues that need to be addressed is illustrated in Fig. 1. The key to this transition is to develop a detailed plan for the transition. To do this, existing plans and agreements must be reviewed and updated according to current agency requirements, priorities and budgets. Part of the trade space in making these decisions is determining the appropriate level of service for each data product with the understanding that this will likely vary from product to product. It will also be necessary for a process to be defined to make the decisions on the products that is transparent and documented and gets appropriate input from each of the involved agencies and the communities that they serve. This process must also work within the existing management processes that have been established by the agencies. NASA and its operational partners, NOAA and the USGS are working together to address these issues and some initial plans and transitions are underway. However, work on the transition of EOS data is still in the earliest stages.




Figure 1. Planning process for transferring data to the LTA

USGS

The USGS has a mandate from the Department of Interior (Land Remote Sensing Policy Act of 1992; Public Law 102-555, 15 U.S. Congress 5601) to establish a public-domain archive of satellite data of the Earth's land surface. The resulting National Satellite Land Remote Sensing Data Archive (NSLRSDA), managed by the USGS EROS Data

Center, includes land and coastal observation data captured by satellites and augmented with data from other government, commercial, and international sources. The major focus of NSLRSDA is to manage and preserve those data so that they remain accessible to a broad range of users over time. This entails employing state-of-the-art transcription and archiving technologies, transferring the data to new media as needed, investigating and implementing advanced media and data storage systems for long-term data preservation, and exploring new ways to make data accessible. A fifteen member Advisory Committee, in accordance with provisions established by the Federal Advisory Committee Act, has also been established to provide advice and council on guidelines and rules relating to NSLRSDA data acceptance, maintenance, preservation, and access management policies. The Advisory Committee meets twice a year (<http://edc.usgs.gov/archive/nslrda/advisory/charter7.html>).

The USGS EROS Data Center utilizes National Archives and Records Administration (NARA) endorsed processes for NSLRSDA and other USGS LTA collections to recommend which data sets are kept or what data may be delivered to NARA in accordance with records management practices and schedules.

The Land Processes (LP) Distributed Active Archive Center (DAAC) was established as an active archive at the USGS's EDC as part of NASA's Earth Observing System (EOS) Data and Information System (EOSDIS) initiative to process, archive, and distribute land-related data collected by EOS sensors, thereby promoting the interdisciplinary study and understanding of the integrated Earth system. The LP DAAC, by design, is a temporary archive until the end of missions, at which time the data is migrated to a long-term archive.

Prior to the launch of Landsat 7, Terra and Aqua, the LP DAAC distributed a variety of prototype and precursor data and products commonly known as Version 0 Data and Products. These data and products have transitioned from the LP DAAC to the USGS long-term archive. The Version 0 products include:

- Advanced Solid-state Array Spectrometer (ASAS)
- Advanced Very High-Resolution Radiometer (AVHRR)
- Global Composites
- Global Normalized Difference Vegetation Index (NDVI) CD
- Orbital Segments
- Aircraft Scanners
- Global Land Cover Characterization (GLCC)
- Global Land Cover Test Sites (GLCTS)
- GTOPO30 and Hydro 1k
- Landsat 7
- NASA Landsat Data Collection (NLDC) MSS/TM
- North American Landscape Characterization (NALC)
- Spaceborne Imaging Radar-C (SIR-C)
- SIR-C Educational CD

In order to maintain consistency and cohesiveness for the science users, the Version 0 data is still searchable and orderable through the [EOS Data Gateway \(EDG\)](#), formerly known as the Version 0 Information Management System (V0 IMS). As the priority dictates, the Version 0 data will be moved from EDG access to an appropriate USGS client for user access.

NOAA

The National Environmental Satellite, Data, and Information Service (NESDIS) is responsible for the collection, archiving, and dissemination of environmental data collected by a variety of *in situ* and remote sensing observing systems operated by the National Oceanic and Atmosphere Administration (NOAA) and by a number of its partners [e.g., National Aeronautics and Space Administration (NASA)]. To prepare for large increases in its data holdings, NESDIS initiated the planning and development for a Comprehensive Large Array-data Stewardship System (CLASS) that provides archive and access services for these data. CLASS must be able to handle the data flow from current satellite-based systems [e.g., Polar-orbiting Operational Environmental Satellite (POES), Geo-stationary Operational Environmental Satellite (GOES), and Defense Meteorological Satellite Program (DMSP)] ground-based systems [e.g., Next Generation Weather Radar (NEXRAD)] and *in situ* systems [e.g., Automated Surface Observing System (ASOS)]. It must also be structured to handle the large increases in data that will come from planned satellite launches [e.g., Meteorological Operational satellites (Metop), National Polar-orbiting Operational Environmental Satellite System (NPOESS), NPOESS Preparatory Project (NPP), and Earth Observing System (EOS) satellites].

CLASS will be operational at two locations, the Office of Satellite Data Processing and Distribution (OSDPD) facility at Suitland, MD and the National Climatic Data Center (NCDC) facility at Asheville, NC. Each facility has similar hardware and identical software, is capable of assuming the overall CLASS load at any given time, and is operational at all times. During normal operations, both facilities share the processing load. The most important difference between the two CLASS facilities is the tape robotic systems and associated Hierarchical Storage Manager (HSM). The interface between CLASS and NCDC's tape robotic systems is defined and described in the CLASS-NCDC Archive and Recall System (NARS) Interface Control Document (ICD).

CLASS provides life cycle capabilities for archiving, distribution, preservation, and operation, such that all approved campaign array-data may be preserved as defined by existing National Archives & Records Administration (NARA) and NESDIS archive policies, distributed as requested to customers, and available for disaster recovery. The scope of these capabilities includes the ability to scale system functionality to continuous growth in campaigns and the preservation needs of the data.

Multiple departments of NOAA contribute to the CLASS design process. The technical management team includes members from the Office of System Development (OSD), National Climatic Data Center (NCDC) and the National Geophysical Data Center (NGDC).

The CLASS Project Management Team (CPMT) must define the detailed policies for lifecycle management. The general organizational objectives are as follows:

- Development of software is geographically distributed, but controlled by a centralized Configuration Control Board.
- Installation of the system will be at two sites, leveraging existing facilities.
- Operations for archive and distribution of data sets is fully automatic and depends upon electronic transfer for source data.
- Operators are able to conduct supporting activities remotely, using a secure interface. Examples of supporting activities are customer help, reconfiguration of hardware for load balancing, and data migration.
- Operators coordinate with the housing facility, but system administration of hardware is the responsibility of the housing facility site management.
- Unless a disaster contingency plan is in progress, each housing facility must have one operational system.

The policies cover the following four categories of capabilities and responsibilities:

Archiving – take in data, catalog, move to archive

Primary data sources are limited to NOAA, NASA, and Department of Defense (DoD) institutions that require long-term preservation and distribution of non-classified data.

The Archive, Access and Distribution System flexibility must be sufficiently robust to cover all the needs of the seven sets of large-array data. Other sources of data (e.g., in-situ) may be added as well to the system, by application to the Data Archive Board and the Board's acceptance.

Any data set must conform to certain requirements on content, format, and handling as defined by agreement between NESDIS and the data provider to be acceptable to the system, but some changes in the system may be required to accommodate each new data set. Based upon direction from the Data Archive Board and/or CLASS Oversight Group, the CPMT initiates analyses and conducts reviews to assess the compatibility of each new data set with the Archive, Access and Distribution System, and determines the system changes required to enable the acceptance of this data set.

Preservation – backup, duplication, migration

The System provides the ability to preserve data from loss. Minimally, preservation requires that two copies of a data be kept at all times. A minimum of 50 miles must separate the storage locations of those copies.

Distribution – Internet-interface, user access, order fulfillment

Automatic distribution of electronically transferable data sets is the preferred method of delivery; the system will be able to provide offline delivery of data. Pricing is not the responsibility of the Archive, Access and Distribution System. The system supports an interface to an Order Management System (OMS) and provides necessary and sufficient information to the OMS so that the OMS may calculate pricing and carry out charging activities.

Operation – people managing the system, policies

The management team publishes and enforces policies for the operations of the resources in normal and contingency modes of operation. The team provides plans to accommodate the continuous growth of data sets and types of data sets that the system must support.

Concluding Remarks

It is clear that the members and advisory committees of the Earth science research and applications communities view NASA's Earth science data as a national resource and emphatically state that it must be preserved and made accessible to support critical science research which may have significant impacts on important policy decisions. They have defined the need for a national infrastructure for LTA of Earth observation data and have recommended that the LTA address the issues of **setting priorities** for the LTA, identifying the **data content** of the LTA (including documentation required for long term use of the data), and specifying the **stewardship and data services** to be provided by the LTA. NASA and its operational partners, NOAA and the USGS, have been working together to follow this guidance.

While concerns over the long-term preservation of data put the focus on the LTA, the studies by the agencies have recognized that this is truly a data lifecycle issue. Plans for the long-term preservation of the data should be addressed during mission formulation and reviewed throughout the data lifecycle to ensure that facilities and budgets to maintain the data are available when they are needed. Also, for the data to be continually useful, complete documentation must be generated and maintained with the data as it transitions over its lifetime. This is sometimes difficult to recognize in the earlier stages of the mission when the science teams and user support specialists of the active archives are available to share their expertise, but it is critical when those teams have disbanded and the data is in the LTA.

The three agencies have been working together to specify preliminary requirements, guidelines and plans for the transition of data from the NASA active archives to the LTA facilities. However, changes in agency directives and priorities, changes in the needs of the communities they serve, and technological advances will always be occurring and necessitate periodic reevaluations of the plans. The greatest challenges will be the definition of the processes that need to be defined to make final decisions related to these transfers, who needs to have a voice in these processes and how these processes work in conjunction with established working groups and procedures that currently govern the work of the three agencies. Each agency's actions are constrained by Administration and Congressional direction, advisory group guidance, priority, policy, and funding. It will take some considerable effort to define a new process that can be compatible with these existing mechanisms.

References

- [1] "Global Change Science Requirements for Long-Term Archiving", Report of the Workshop, Oct 28-30, 1998, U.S. Global Change Research Program [USGCRP] Program Office, March 1999.
- [2] "Ensuring the Climate Record from the NPP and NPOESS Meteorological Satellites", National Research Council, Committee on Earth Studies [NAS-CES] 2000).
- [3] "Workshop on Long-term Archiving of EOS Data", January 29-30, 2002, EOS Science Working Group on Data, unpublished.

Session Future Prospects

DIGITAL LIBRARY: IMPROVING THE ACCESSIBILITY OF THE RUSSIAN SATELLITE DATA IN SUPPORT OF THE ENVIRONMENTAL MONITORING

Efim Kudashev

Space Research Institute of Russian Academy of Sciences

84/32 Profsovnaya str., Moscow, 117997 Russia

Email:eco@iki.rssi.ru

1. RUSSIAN REMOTE SENSING SATELLITE DATA

Today, Digital Libraries of satellite data are being created in several countries in support of Earth investigation programmes. These national information systems use satellite data flows to deal with various problems of Remote Sensing. For many years Russia has been carrying out an extensive programme of Earth Observation using a range of satellite remote sensing instruments. Russian satellites have collected large volumes of data over the years, and exploitation of this data world-wide could have a wide variety of uses, notably environmental monitoring and natural disaster mitigation [1,6,7]. Due to activities such as IGOS (Integrated Global Observing Strategy) and DMSP (Disaster Manager) there is a need for a global infrastructure of data. GMES (Global Monitoring for Environment and Security) aims to link major satellite and surface based systems for global environmental observations of the atmosphere, oceans, land and biota. For Russian data to be considered as part of this, it needs to be accessible, linked to, and interoperable with, other national and international information systems. However, despite the large volumes of data generated by these systems they are not exploited as much as they could be, especially outside Russia. It is important that such potentially useful datasets are utilised to the maximum extent possible in order to achieve maximum cost effectiveness from the previous space investments and to justify future expenditure on new and continuing programmes. The INTAS Project IRIS (Integration of Russian Satellite data information resources with the global network of Earth Observation Information Systems (IRIS) activity is linked to ESA'S INFEO (Information on Earth Observations) and is aimed at increasing international cooperation. Access to satellite datasets archived at the IRIS is achieved through INFEO Data Gateway <http://iris.iki.rssi.ru>. The main aim of IRIS project is to improve the accessibility of Russian satellite remote sensing data, and to make recommendations for the further development of regional satellite archives in Russia. A number of objectives have been defined, including generating Russian satellite information resources, linking with the European INFEO (metadata/ catalogue) system in order to provide interoperable catalogue access and ensuring that end users are involved.

2. INFORMATION TECHNOLOGIES AND INFORMATION RESOURCES FOR SPACE ENVIRONMENTAL MONITORING

Developments in GIS and WWW technologies are providing new tools for finding and retrieving data over computer networks. Efficient organization of Remote Sensing Archives and open access to Data Bases are founded on the Web technology (access to data including data search and request). The development of satellite natural-resource information software targets the following problems: Real Time Earth Observation, thematic processing of Remote Sensing Data and filling in of the Digital Archive, ecological monitoring of the environment, ecosystem condition evaluation through space techniques, access to hydrometeorological data from around the globe.

Digital Libraries of Remote Sensing use satellite data flows to deal with various problems of Remote Sensing posed by both scientific community and specific industrial needs:

- The structure and generation of Russian satellite information resources (including database generation, archiving, metadata generation and catalogue creation).
- Linking regional satellite archives for acquisition, processing and distribution of remote sensing space data in a consistent manner to provide interoperable catalogue access and data delivery over the Internet.
- Integrating Russian satellite data information resources with the European INFEO system and other international networks of Earth Observation information systems (e.g. EOSDIS in the USA and EOIS in Japan).
- Ensuring end-users of the satellite data are involved in the project to ensure that the research is well directed and addresses the issues of concern to data users, particularly in the fields of environmental monitoring and natural disaster mitigation.
- Linking existing information dissemination and teaching activities to ensure efficient propagation of the results to a new generation of scientists.

There are 7 partner organisations involved in the project, 1 from the UK, 1 from Italy and 5 from Russia.

The Partners involved in the project were: IRE (Institute of Radio Engineering and Electronics, Moscow) <http://www.ire.rssi.ru/cpssi/cpssi.htm>, IKI (Space Research Institute, Moscow) <http://iris.iki.rssi.ru>, Forest Institute (Krasnoyarsk) <http://www.krasn.ru/Forest>, IAPU (Institute Automation and Control Processes, Vladivostok) <http://www.dvo.ru>, Faculty of Geography, MSU (Moscow State University). Coordinator: Space QinetiQ (UK) <http://www.space.qinetiq.com>. TERMA Electronica (Italy).

The IRIS is a consortium of researchers working together to develop distributed Digital Library for heterogeneous Earth Science Satellite Data. The tasks of IRIS Project are:

- Improvement of the infrastructure of archive and data exchange systems for processing and distribution of optical, IR and microwave information.
- To identify shortcomings in current technologies and provide inputs into the ongoing work of standards and co-ordinating bodies such as ISO TC211, CCSDS Panel2 and CEOS (Committee on Earth Observation Satellites).
- To create new information databases and information services using the experience databases of IKI, IRE, IAPU and Institute Forest, of covering unique multi-year (1974-1994) optical and microwave aircraft and satellite measurements in North-West Pacific (IKI) together with similar data collected by IRE over the territory of the former Soviet Union and Europe.
- To determine how best to provide users with prompt ecological and meteorological information available for environmental protection.
- To investigate the forecast methods for natural and technogenesis phenomena based on optical, IR and microwave information from RESURS and METEOR satellites, in particular, forecast methods for large-scale atmospheric phenomena (e.g. hurricanes).
- To determine how best to integrate Russian archives to create a Digital Library based on the concept of an interoperable distributed archives.

3. FROM DATA BASES TO DIGITAL LIBRARY OF REMOTE SENSING DATA

With item “improve the accessibility of Russian satellite remote sensing data”, it was recognized at the start of IRIS project that developments in Geographic Information Systems (GIS) and World-Wide Web (WWW) technologies were providing new tools for finding and retrieving data over computer networks. However, it was not clear at that time how best to exploit these new tools in the field of satellite Earth Observation in the Russia.

The specific objectives of the project in this area were to investigate issues concerning i) the structure and generation of Russian satellite information resources including database generation, archiving, metadata generation and catalogue creation; ii) catalogue access and data delivery over the internet; iii) integrating Russian satellite data information resources with the European INFEO and global network of Earth Observation information systems. All these objectives were pursued and successful results obtained. In particular, some new catalogue databases were created and existing catalogue systems were reviewed and updated to conform to international standards. Web-page servers have been established to allow these catalogues to be searched through bi-lingual interfaces. High-level dataset descriptions were created and submitted to the International Directory Network (IDN). Data users are now able to search the IDN and become aware of the existence to the Russian datasets. Once discovered, the users can then access the specific websites to carry out more detailed searches and order the data required. In addition to the general cataloguing and establishing web access, two of the partners have installed “INFEO gateways”. This software package allows the associated catalogues to be connected to the international INFEO catalogue searching system. Data users are now able, through a single entry point, to enter searches to all the associated catalogues around the world and see search results from Russian catalogues if they satisfy the search criteria. The success of this work has been reported through papers in scientific journals and at conferences [2-5]. However, the true success is not achieved through advertising but through users actually finding the data through searching, without knowing if it existed beforehand. It is also planned to report the activities to the next meeting of the International Committee on Earth Observation Satellites (CEOS), Working Group on Information Systems and Services (WGISS). This body was the driving force behind the original development of the INFEO system.

With item “support the exploitation of such data”, a range of processing algorithms have been developed and new applications pursued. A number of new data sources have also been established, including digitised data from old satellite photographic images and data from Chinese satellites. The focus of the data exploitation has been on ocean data from the Far East marginal seas and on forest management in Siberia.




Fig. 1. Model of Individual Data Archive.

The IRIS project has allowed firm relationships to be developed between the partners and future interaction will continue after the end of the project. The initial goal of connecting Russian satellite catalogues to the international INFEO system has been achieved although the work will need to continue with the addition of new datasets as they come along and upgrades to the software will need to be installed as computers and associated software continue to improve. However, the initial barriers have been overcome and it is expected that the on-going maintenance can now be achieved on the back of new data exploitation projects as they are established.

IKI have shown how to integrate regional satellite datasets and other archives to create a single EO Information System for sustainable development based on the concept of an interoperable distributed archive. New learning materials, published on the IKI Web server, have contributed to EO education. IAPU report that the results achieved thus far would impact on the fields of typhoon diagnosis, noise filtration of geostationary satellite data and also the automatic detection of sea surface eddies using satellite data. The results obtained by Institute Forest have more of an environmental impact, as the potential capability of satellites to detect disturbances such as pollution, gold mining and pest impact has been shown. This capability can be used on natural and anthropogenic disasters. The Internet-based seminars created, and hosted, by MSU have impacted the level of education in Russian universities. The atlas of space imaging highlighted other uses of remote sensing to a wider community. Training in GPS and other field equipment, and preparing of a text-book (Practicum), support the introduction of Russian specialists into a modern informational society [4].

4. EUROPEAN INFEO (INFORMATION ON EARTH OBSERVATIONS) SYSTEM

The INFEO system offers access to Earth Observation information, services and data catalogues around the world. User easily exchanges data with the INFEO system. This simple point-and-click, search and order WWW interface is based on a hierarchical organization of data, displayed as tables, and will always return non-empty results. An archive (See Fig. 1) can hold various types of data ranging from satellite images and products processed from the images, to observation data and statistics. An archive may also contain documentation. Digital library is the totality of many individual data sets, collected and held separately by many different organizations [2,3]. The INFEO system offers access to Earth Observation information, services and data catalogues around the world.




Fig. 2. The architecture of a CIP/ODBC Gateway.




Fig. 3. CIP/ODBC Gateway.

The INFEO system comprises Middleware Nodes (MWNDs) that, taken together (see Fig. 3), define a distributed hierarchy of collections of data [5]. Figure 3 shows a schematic INFEO architecture of several MWNDs each of which conceptually owning a number of gateways (GWs) or servers. The INFEO CIP/ODBC Gateway is one example of these components. The various databases and servers shown are typically legacy provider systems, e.g., existing archives and inventories. All Users can receive the access to INFEO data over MWNDs. A MWND (middleware node) in a top-level component of INFEO.

CIP/ODBC Gateway consists of following components: Retrieval Manager; Translator.

Figure 2 shows the architecture of a CIP/ODBC Gateway illustrating the management of user-driven CIP sessions and translation of CIP into ODBC. The installation of the CIP/ODBC Gateway and the configuration of its RDBMS Translator by mapping CIP to the schema of the relational database.

CIP/ODBC Gateway is intended to provide access to legacy databases by translating CIP requests from a MWND into ODBC (i.e., SQL) suitable for execution on the database. The Gateway is an example of an ICS Retrieval Manager. A Retrieval Manager (See Fig. 4) manages CIP/Z39.50 sessions and services (and may be installed at) each catalogue site as well as each search server. It is used to integrate together the local catalogue systems and provide communication between users and other catalogue site Retrieval Managers.

Pre-Packaged Retrieval Manager consists of following components:

Customer User Interface (CUI); ICS Retrieval Manager (RM, Translator and Data Manager); Databases; Associated software.




Fig. 4. Pre-Packaged Retrieval Manager Architecture.

5. EARLY STAGE OF CREATING A DIGITAL LIBRARY

Already at the early stage of creating a Digital Library of Remote Sensing data, it is essential to work out its adequate structure necessary for efficient data retrieval from the archive. The Digital Library structure is elaborated based on the understanding of typical requests of potential archive users. The experience of functioning space archives shows that user requests primarily focus on data representation levels, the name of the project under which the data is obtained and the name of the sensor that provided the data. Therefore, the archive should be divided into segments corresponding to different data representation (process) levels with each segment subdivided into data sets related to a certain project and instrument (sensor). Efficient organization of information resources and open access to spatially distributed experiment data are founded on the information services of Internet, i.e. on the Web technology. Appropriate metadata management systems are built to provide for the collection and distribution of experiment data and thematic processing results; while the Digital Library is linked to the regional archives of environmental monitoring via Internet. An important element is the elaboration of interface, archiving and network data exchange structures. This calls for the development of search engines and a remote interactive access regime for external users via Internet to catalogues of experiment data and processing results and the realization of the on-line access mode. Gateway software installation, and Gateway software configuration are presented. Digital Library uses a satellite image dataset model and CIP – Catalogue Interoperability Protocol.

Gateway software installation

FTP-access give us access to Gateway soft ODBC package (Open DataBase Connectivity) – Open Interface for interaction to database (See Fig. 5). Microsoft standard regulates the improved access of the applications to existing databases, allowing users without specific knowledge of data to search science data holding, retrieve high-level description of data sets and detailed descriptions of the data inventory, view browse images.




Fig. 5. Schematic of INFEO Gateway.




Fig. 6. The Concept of a Collection.

Collection can be form: 1) the product descriptors or 2) collection descriptors (See Fig. 6). User can search for and order data. Search Types: i) detailed descriptions about data sets, data sources, instruments, projects; ii) collections of observations of data (granules).

Results of Search are improved online access to existing data.

Schematized Collection Model includes: i)Terminal Collection (with numbers '1.x') – Earth Observation product data only; ii) Non-terminal Collection.




Fig. 7. Protocol Z39.50 Service Model.

A distributed hierarchy of collections of data

We will describe in this section the hierarchical approach. The definition of collection to be considered in designing a browsing and navigation service are: a grouping of item descriptors that have commonality. Within a single grouping or 'collection' the item descriptors must be of the same type (i.e. the items can be described with the same set of attributes). In the Earth Observation Collections at a MWND only contain other collections. No product or document descriptors are held at the MWND. Each collection is defined by a descriptor that describes the data it contains by means of metadata and typically represents a theme, e.g., a collection of data from a particular satellite or sensor processed in a particular way. Each terminal collection in the hierarchy represents the data available at a particular provider gateway or server. The intention is that users find collections of data of interest based on the collection metadata (by performing a collection search at a MWND) or by perusing the collection hierarchy and then finds data of interest within particular collections (by performing a product search).

Catalogue Interoperability Protocol (CIP)

Standard of protocol CIP is based on Z39.50. Protocol CIP is a profile of Z39.50. Z39.50 Protocol means for Search and Remote Sensing data presentation [6]. In the case of Digital Library CIP uses Z39.50 means for distribution search (See Fig.7).




Fig. 8. INFEO Simple Search.

6. CONCLUSIONS

The paper describes technological (software) aspects of building distributed Digital Library for Satellite Based Data. A Satellite Image Datasets is open for researchers and can be accessed via Internet. Therefore, we are building the Digital Library as Information System with increasing the number of Earth Observation Data from now on to be further easy to utilize. The example of INFEO Simple Search to get Satellite Based Data, including Data Exchange among satellite archives showed on Fig. 8.

Those successful demonstrations and uses Digital Library of IRIS systems how it is feasible for the data interoperability based on Open GIS Consortium (OGC) at <http://www.opengis.org>, Web Coverage Service (WCS) at <http://ip.opengis.org/ows1/docs/01-018r1.doc>. It also demonstrates that providing interoperable, personalized, on-demand data access and services to the data users greatly enhances the use of Remote Sensing Earth Observation data in scientific research and applications.

REFERENCES

- [1] Kudashev E.B, Kravtsov Yu, and Golomolzin V.V., "The Satellite Techniques in Early Warning Systems," in *Early Warning Systems for Large Towns and Megacities*, Chapter 8.2. J.Zschau and A.Kuppers, Eds. Berlin, Heidelberg: Springer-Verlag, pp. 741-744, 2003.
- [2] Kudashev E., "From Data Bases to Digital Library of Remote Sensing Data", in *Proc. 23 EARSeL Symposium*. Abstract Book. pp.142-143. Ed. Univ. of Ghent. 2003.
- [3] Kudashev E.B., Balashov A.D."Integration of Digital Library of Remote Sensing Data to Earth Observation Information System", in *Proc. Fifth Nation. Russ. Res. Conference. Digital Libraries: Advanced Methods and Technologies, Digital Collections*, pp.88-95. S.-Petersburg Univ.Published Press, 2003.
- [4] Wyn Cudlip, Efim Kudashev, Dmitry Kurtener, "Conception of integration of Russian satellite data information resources with the global network of earth observation information systems (IRIS)", in *Agrophysical and Ecological Problems of Agriculture in the 21st Century*. Vol. 3. P.127-133. SPBISTRO, St. Petersburg, 2002.
- [5] Kudashev, E.B., Yu.F. Knizhnikov, V.I. Kravtsova, V.P. Myasnikov, O.V. Tutubalina, "INTAS Project: Remote Sensing Internet Technologies based teaching", in *Proc. 2002 Intern. Geoscience and Remote Sensing Symposium IGARSS'02*. V. 1. pp. 539- 541. IGARSS:Toronto, Canada, 2002.
- [6] "CIP Search-B. Software User Manual". Space Applications Inst. GB, 2000.
- [7] Kudashev E.B., Kravtsov Yu.A., Kharuk V.I., Myasnikov V.P., "Using the Russian satellite data in support of environmental monitoring and emergency forecast for vast syberian and from Moscow to Arctic territories", in *21st EARSeL Symposium. Observing our Environment from Space/ New Solutions for a New Millennium*, p. 83. European Association of Remote Sensing Laboratories. Paris. 2001.
- [8] Kudashev E.B., Kravtsov Yu.A., Myasnikov V.P., "Remote Sensing and development of databases/ large GIS for Environmental Monitoring of the Moscow megacity", in *Proc. IEEE/ ISPRS Workshop URBAN*. Univ. "La Sapienza". Rome. 2001.

Providing Authentic Long-term Archival Access to Complex Relational Data

Stephan Heuscher^a, Stephan Järman^b, Peter Keller-Marxer^c, and Frank Möhle^d

Swiss Federal Archives, Archivstrasse 24, CH-3003 Bern, Switzerland

^astephan@heuscher.ch, ^bstephan.jaermann@bar.admin.ch, ^cpeter.keller@acm.org, ^dfrank.moehle@bar.admin.ch

We discuss long-term preservation of and access to relational databases. The focus is on national archives and science data archives which have to ingest and integrate data from a broad spectrum of vendor-specific relational database management systems (RDBMS). Furthermore, we present our solution SIARD which analyzes and extracts data and data logic from almost any RDBMS. It enables, to a reasonable level of authenticity, complete detachment of databases from their vendor-specific environment. The user can add archival descriptive metadata according to a customizable schema. A SIARD database archive integrates data, data logic, technical metadata, and archival descriptive information in one archival information package, independent of any specific software and hardware, based upon plain text files and the standardized languages SQL and XML. For usage purposes, a SIARD archive can be reloaded into any current or future RDBMS which supports standard SQL. In addition, SIARD contains a client that enables ‘on demand’ reload of archives into a target RDBMS, and multi-user remote access for querying and browsing the data together with its technical and descriptive metadata in one graphical user interface.

The urgency of deep-infrastructure solutions for long-term digital preservation and archiving was clearly formulated in the late 90’s of the past century by national archives and libraries [1–7] as well as space agencies and institutions in earth observation, oceanography, and astronomy [8, 9]. The finding of problem statements and strategies is still in progress and was recently supported by a charta of the United Nations [10]. Deep and lasting solutions are still not available, and for national archives and libraries (faced with very heterogeneous types of content) it is broadly accepted that digital collections are growing at a rate that outpaces their ability to manage and preserve them [11].

There is an ongoing process of recognition that long-term digital preservation poses similar problems in diverse disciplines. Despite of different vocabularies used by different communities, research and development in the field can only be successful in a joint effort. During the past decade, however, decisive progress was achieved in analytical and conceptual work [12–14]. Furthermore, the Open Archival Information System (OAIS) reference model became widely accepted in diverse disciplines and covers a full range of archival information preservation functions including ingest, archival storage, data management, access, and dissemination. It has become the international standard ISO 14721:2003 [15] and may be applicable to any archive as it does not refer to specific implementations or archival strategies but merely provides a common terminology and functional framework to discuss and assess different implementation approaches.

Many current development projects focus on the archiving of digital images (digitized photographs or paper documents) and sound recordings, while more complex types of digital information have been neglected, even though their relevance for governmental and scientific activities has drastically increased during the past decade. In particular, a recent international workshop on the long-term preservation of databases [16] revealed that many archives do have a long-standing practice and experience in ingesting and preserving relational data, but their daily work is constrained to the treatment of rather simply structured data sets, requires extensive manual work by archives personnel, and does not allow for smooth and standardized integration of data and descriptive meta data on a level required to ingest, preserve, and provide access to complex relational databases.

In this paper we present a method and application named “Software-Invariant Archiving of Relational Databases” (SIARD), developed at the Swiss Federal Archives. It completely detaches typed relational data from almost any relational database management system, while still retaining most of the original data logic and integrating data and metadata in one archival information package that is based on text files and standardized technologies. In Section II we discuss the technical and intellectual complexity of relational data in modern database systems, the resulting problems for long-term preservation, and its relevance to archives. The objectives in the development of SIARD are described in Section III, while Section IV covers SIARD’s system architecture, workflow, features, and development platform.

I. INTRODUCTION

Relational data is one of the oldest forms of structured information representation, intuitively used already centuries before the “digital age”. With the rise of computer technology, the introduction of mathematical formulations of the relational data model in the mid-20th century, and the international standardization of

a corresponding data definition and query language, relational data has become an omnipresent method to organize data for electronic data processing in almost every field of work, from business activities to scientific research and government administration.

During the past two decades, usage has developed from processing single-table data files with specific application software to generic relational database management systems (RDBMS). These have internal mechanisms for logical and physical organization of arbitrary relational data models, are able to physically store terabytes of data, cover rich data types (including internal procedural code), enable multi-user transactions, and provide internal data life-cycle management. The definition, representation, management, and query of relational data was thereby standardized and separated from specific application logic and application software that operates on the database. As a result, RDBMS have become core components of almost any type of digital information system.

It is obvious that this development has decisive impacts on the work of those institutions which are charged to collect or accept digital data from various data sources, to make it broadly accessible, and to preserve it over decades: national archives and libraries, science data archives, or business companies being under special legal regulations for long-term data retention (like, for example, the pharmaceutical sector [19]).

II. COMPLEXITY AND RELEVANCE

A. Technical Complexity of Relational Data

One consequence of the progress in database technology is that relational data and relational databases become highly complex. They often consist of hundreds of linked tables (i.e. physical representations of relational entities¹), which makes it impossible to handle and query table data outside an RDBMS if these links become broken or cannot be managed automatically anymore.

Furthermore, any data item in a RDBMS has a precisely defined data type and domain. (Entire tables or integral parts of them may likewise have types.) Apart from basic data types (for example integer and real numbers, dates, and character strings), low-level types like large binary objects, complex and inheritable user-defined types, and multi-lingual character encodings are widely used in modern databases. As for the linking of tables, the connection between the data and its data types and domain definitions is not preserved when table data is trivially exported to external plain text files.

In addition to those entities and features mentioned above, modern RDBMS include, for example:

- Check constraints and assertions: For a single column of a table or a set of entire tables, assure that changing or entering data does never violate defined data types, quantitative restrictions, or value domains. In particular, it can be assured that data items in a table column will never be empty.
- Views: Assemble selected parts of several tables and operate on them as customized, virtual tables.
- Triggers: Force the RDBMS to initiate timed operations on data when user-specified conditions are met, for example log and audit user activities.
- Functions (basic and user-defined): Perform numerical calculations, conversions, or character operations on data items or sets.
- Stored Procedures: Store and execute programs inside the RDBMS to perform common or critical tasks which are not part of the specific application software outside the RDBMS.
- Foreign Keys: Ensure referential data integrity, i.e., automatically prevent that values can be stored in rows of one table if there are no corresponding values in referencing entries within the database.
- Grants and Roles: Define user profiles and assign or withdraw privileges, for example to create new tables or access certain parts of the database.

¹ We will not discuss the relational model in this article [17, 18]. It will be sufficient to think of *tables* which consist of one or more table *columns* and one or more table *rows*. The points of intersection of columns and rows contain *data items* which have a *data value* and a *data type*. If (and only if) there never occur duplicate rows in a table, then the table is a *relation*, and its columns and rows are also called *attributes* and *tuples* (or *records*), respectively. Using relational algebra and calculus, several relational tables can be managed and manipulated jointly.

For long-term preservation in national archives, relational data is collected from many different database systems and has to be retained and kept processable and accessible for decades. It is therefore essential to store and maintain the databases independent of any specific and short-lived products (or at least transferring them all into only one preferred product). In fact, most RDBMS use the same language for the definition of the internal logical organization of data, namely the declarative (i.e. non-procedural) Structured Query Language (SQL). Despite of its name, the scope of SQL also includes the definition of data structure and the manipulative operations on data stored in that structure [17].

The development of SQL started in the 1970's, leading to the international standard ISO/IEC 9075 in 1987, and evolved in four main stages through SQL-89, SQL-92, SQL:1999, and recently SQL:2003 [20], while the size of the standard has grown from 120 to over 2'000 pages. SQL:2003 and SQL:1999 are fully upward compatible with SQL-92. The standard language keywords are structured in three subsets: Data Definition (DDL), Data Manipulation (DML), and Data Control and user authorization (DCL). To increase acceptance by vendors, the standard defines three levels of conformance and implementation: entry, intermediate, and full level. The mandatory part of SQL:1999 and later is called the "Core" of SQL and described in Part 2 (Foundation) and Part 11 (Schemata) of the standard.

Aside from revisions to all parts of SQL:1999 (e.g. new data types and functions that return entire tables [21]), the 2003 edition contains the new part: "SQL/XML" defines a minimal handling and integration of text-based data structured by the Extensible Markup Language XML [22]. This includes [23] mappings between tables and XML documents, SQL data types and XML Schema [24] data types, and RDBMS implementation-specific character sets to Unicode [25–27] character encodings. Additionally, the related standard "SQL Multimedia and Application Packages" ISO/IEC 13249:2003 [28] defines a number of packages of generic data types common to various kinds of data used in multimedia and application areas, to enable storage and manipulation of such data in a relational database.

1. *Standardized – Really?*

SQL is an internationally standardized and comprehensive language for the definition, description, query, and manipulation of relational data and databases. It is widely used since almost 25 years, developed upward compatible, and will probably play a key role for another 25 years. Since most RDBMS are based on SQL (and most vendors claim compliance with the standard) one could assume that relational database definitions are independent of any specific RDBMS product. Unfortunately, this is far from being true. In contrast to standardized programming languages like ISO-C or ANSI-Fortran, SQL-based database layouts and SQL code can rarely be ported between different RDBMS without major modifications and loss of functionality.

There are two main reasons for severe incompatibilities. First of all: Although the SQL standard today comprises over 2'000 pages, it is far from being fully self-contained. In contrast, SQL:1999 explicitly identifies 381 so-called implementation-defined items and 137 so-called implementation-dependent items [17]. Their implementation is left open for any manufacturer of RDBMS products. As long as a manufacturer completely documents all implementation-defined items, the product can rightly claim to comply with the SQL standard, though it differs from all competing products. (The precision of the SQL integer data type is a simple example of an implementation-defined item.)

The second reason is that most of today's RDBMS implement only (and sometimes faultily) the core and the entry level of the standard completely, but add plenty of non-standard, product-specific enhancements, leading to different "flavors of SQL". These include [29] new additions to or modifications of, for example, data types, functions, operators, behavior and syntax of SQL statements. Additionally, almost all RDBMS products use their own procedural programming languages for stored routines (Oracle PL/SQL, Postgres PL/pqSQL, Microsoft T-SQL, PL/Perl etc.) rather than implementing the standard's procedural language SQL/PSM.

Finally, it is an almost trivial remark that modern RDBMS move the physical storage of the data from the operating system (file level) to the application level (the internal storage of the RDBMS).

2. *"SQL for Archiving"?*

Considering the imponderables discussed in the previous subsection we can draw the following conclusion with respect to archival institutions: If they have to collect and ingest relational data from various database management products for the purpose of integrating them, to preserve them over long periods of time, and to make them broadly accessible, then these institutions are faced with a Sisyphean task: The data they have to preserve for long is locked up in short-lived obsolescent and complex software products from a vast diversity of

manufacturers, making the longevity of the data heavily depend on the availability of the products and their versions, the support by vendors, and the existence of the manufacturers. One lesson learned is that long-term preservation of relational databases is much more than just making backups of export or dump files from database management applications, though IT professionals usually use 'archiving' and 'backup' as synonyms.

It is a tempting idea to demand a "SQL for Archiving", defined as a subset of the ISO-9075-SQL standard. It would leave database designers the choice to restrict their databases to layouts suitable for long-term preservation (and transfer of databases among competitive RDBMS products, of course). A similar effort is undertaken by government agencies and industry representatives to define an ISO standard "PDF for Archiving, PDF/A" [30] based on the popular Portable Document Format (PDF) specification 1.4 [31] by Adobe Systems Inc. and the ISO 15930 standard PDF/X [32]. However, the current ISO-9075-SQL standard is probably not suited for a similar attempt: too many important items are left open as implementation-defined and would require adding deep descriptions of implementation details in the standard.

It should be mentioned that since its first edition, ISO-9075 requires SQL implementations to provide a feature called "SQL-Flagger" (feature F812 in SQL:1999) which is able to identify and signal certain kinds of non-standard SQL language extensions used in the specific implementation. The feature is implementation-dependent and its intention is to assist SQL programmers in producing SQL language that is portable among different conforming SQL-implementations. Standard SQL flagging is only required for the entry level of SQL. In fact, most major RDBMS manufacturers have implemented SQL-Flaggers for SQL-92 in their products. The reason is that the U.S. Government implicitly requires conformance with the entry level of SQL-92 for all SQL products in federal procurements. (Conformance with higher levels may be specified explicitly.) The minimum requirements for conformance with entry-level SQL-92 as well as specific features of the SQL-Flagger are specified in the Federal Information Processing Standard (FIPS) 127 [33].

The purpose of FIPS SQL is "to promote portability and interoperability of database application programs, to facilitate maintenance of database systems among heterogeneous data processing environments, and to allow for the efficient exchange of programmers among different data management projects" [33]. Although not explicitly mentioned in FIPS 127, its intention also supports long-term preservation of databases. The U.S. National Institute of Standards and Testing (NIST) used to validate SQL implementations to conform with FIPS-127. Although this procurement specification is still in force nowadays, it was not updated to the 1999 and 2003 editions of SQL, and NIST ceased its product validation in 1997. (A suite of automated validation tests for SQL-92 implementations is still freely available from NIST [33].) To our knowledge, none of the major RDBMS products include SQL-Flaggers for SQL:1999 or SQL:2003 as it would be required by the respective standard editions. Our application SIARD (described in Section IV) has its own, built-in SQL validator. Aside from this, we are aware of only one (third party and commercial) tool which provides such functionality [34].

We conclude from our work, that standard (i.e. generic) SQL may be reasonably exploited for long-term preservation purposes only when data and data logic is actively extracted from database management systems by specialized ingest tools which map different "SQL flavors" to generic SQL, and transparently trace and document those parts which cannot be mapped.

B. Intellectual Complexity and Access

Aside from the technical aspects discussed so far, a successful long-term preservation of databases is only possible if the intelligibility and comprehensibility of the database and its data can be preserved as well. To keep data understandable and meaningful it is indispensable to collect enough technical as well as non-technical metadata and handle it as an integral part of a database archive. Otherwise, there will be a rare chance to understand the meaning and value of the database's content decades after it was archived.

However, most of the meta-data necessary to enable long-term *intellectual* access to the data is not deposited in the database but is provided to the archives on separate (and often paper-bound and hard to grasp) documents. On the one hand, this includes precise and complete data dictionaries, code lists, narrative descriptions of the naming, meaning and usage of single database objects. On the other hand, it is descriptive and archival metadata about the context, creation, purpose, usage, or chain of custody of the original database and the RDBMS used to operate it.

The problem of descriptive metadata that supports intellectual accessibility has to be considered in a broader context of trusted digital repositories [35]: Any serious long-term preservation strategy for any kind of digital content aims to guarantee continuous

- **Integrity:** protection of the data from unintended and intended harm;
- **Intelligibility:** understandability and comprehensibility of the data;

- **Authenticity:** authentication (of authorship and provenance) and reliability; (of evidence)
- **Originality:** data structure and functionality “as close to the original as possible”;
- **Accessibility:** technical readability and usability.

However, due to overall technical obsolescence in a digital environment, these are competitive and conflicting goals. Each archival institution will therefore have to establish its own measures and priorities among the five above-mentioned criteria. The measures will be primarily ruled by metadata, and many authors address the interplay between obsolescent technical infrastructures and continuing guidance of metadata [6, 35–37].

Moreover, since long-term preservation is a costly and laborious task, effective appraisal methods are increasingly important instruments. The value of information and evidence contained in databases is often determinable only through the purpose, design, and context of usage of the original database application. Again, these criteria are measured by means of metadata provided by the data producers.

Another level of intellectual complexity is introduced by an increasing amount of interlinked, federated and temporal database systems which makes it difficult to determine the correct spatial and temporal scope for extracting data for long-term preservation: A database that is selected for archiving may refer to time-dependent master data in another database, or a database does not overwrite or delete any outdated data but rather records all data modifications, using timestamps as multiple primary key components (e.g. valid-time state tables, tracking logs, backlogs, etc.). It will be a challenge for archives to keep such spatial and chronological dependencies and interrelations understandable and traceable, particularly across sequential accessions from the same database.

C. Current practices

Relational data kept by national archives usually consists of plain text files that contain tabulated data with fixed-length or delimited columns [16]. Very often, these files are only derivatives of the original database, produced by denormalization of the original data model to reduce the number of single tables. The data files are usually accompanied by paper documents or microfiches that provide data dictionaries and other descriptive information necessary to understand the content, provenance, and context of the data. At the time when the data was originally transferred to the archives, it has been scrutinized to reveal any inconsistency with the paper-bound documentation.

In typical archival environments, the descriptive documents are kept separately in cardboard boxes which are stored in air-conditioned shelf vaults, while the electronic data files are stored on labeled magnetic tapes or cartridges. Tapes and cartridges are recopied every 5 to 10 years to prevent data loss from degradation of the magnetic media by physical and chemical processes.

One reason for the deficiency in providing long-term physical and intellectual access to data from archived databases is a high heterogeneity of data formats and the fragmentation of data sets into isolated data files, which are hard to handle in bulk. But primarily it is the habitual bipartite treatment of digital data in most archives: While the data itself is accessible and processable by electronic systems, the metadata necessary to understand the data’s content, provenance, and context, is bound to paper documents and often incomplete, erroneous, outdated, not standardized, and hard to grasp. That’s why even a small data collection requires cumbersome and expensive manual work to overcome this divide between electronic and paper documents.

It can be easily concluded from the discussion of the previous sections that the situation will become even more critical in the future: The advance in information technology puts more and more obstacles on the path that digital information has to pass on its way from producers to the archives. Proprietary, closed data formats and technologies quickly become obsolete, and heterogeneous and complex data structures further impede comprehensive data integration in the archives.

Today’s archives become aware that digital information in their custody becomes increasingly volatile rather than persistent. This is particularly true for relational databases since the development of new preservation techniques has mainly focused on other forms of digital content such as images or sound recordings which are easier to handle and more attractive to a broad audience. In contrast, methods for the long-term preservation of data from modern database management systems with increasing complexity and data in the order of terabytes has been neglected.

D. Relevance to National Archives

Relational data is probably the oldest and most widespread type of information among the electronic digital holdings of national archives, typically dating back to the 70's and 80's of the past century. (In fact, even the oldest collections are of no age compared to usual archival time scales.) These collections comprise data from almost any field of activity of governmental institutions and are thus of high information and evidential value, and are increasingly important to researchers in diverse fields such as History, Sociology, Politics, Economics, Meteorology, or Geography.

Nevertheless, national archives nowadays are hardly ever able to provide broad access to their digital data sets on a level of usage comparable to paper-bound or analog electronic holdings. (A few archives have put a lot of effort into it, though, and provide public access to selected collections through the world-wide web [51].) Most national archives are still far from enabling the public to locate government information regardless of format.

In terms of archival appraisal, databases may serve various purposes which are not in the primary focus of national archives to preserve evidence in business records of the federal administration. (Business records are evidence of what an agency has done or decided.) In the opinion of many archivists, databases are mainly used by the administration to store and manage registers, master data records, and statistical (e.g. census) data, thus may be at best considered as finding aids. This may be true for databases of the past.

Nowadays, modern database systems are widely used for managing and recording business processes and transactions. Thus they have become integral parts of almost any record and document management system or E-government web site, and thus often contain high evidential value. It has become an essential necessity for archives to assess the evidential value of databases kept in RDBMS, and to develop appropriate criteria and guidelines for such assessments [38, 39].

Criteria and guidelines are also required for appropriate archival description of relational databases since they are rarely integrated into any filing plan of the agency. In consequence, there is only a purely technical ordering of database records which will not correspond to the logical ordering of business records (e.g. belonging to a business case, dossier, or document). A business record may comprise several technical table records and may extend across several tables, and a single database table is often not a meaningful entity for archival description. It can be a challenge to identify and precisely describe something like a dossier, business record, or document in a relational database made up of dozens of tables [38, 40].

E. Relevance to Natural Science Data and Research

It is broadly accepted that experimental scientific data has often a high long-term value [2, 3, 41–43], and the urgency of solutions for their long-term preservation has been pointed out many times during the past decade [44, 45]. By contrast with national archives, science data archives focus on the information value, while evidential value of scientific data is rarely considered. Moreover, there are completely different criteria and schemes for appraisal and description of scientific data [46, 47].

Scientific data may have a high long-term value because, for example, it cannot be reproduced (e.g. climate and oceanographic data) or it was produced at enormous costs (e.g. high-energy physics experiments or space flight missions). As theories evolve and new questions arise, archived data may be reconsidered and re-evaluated in future research, and may turn out to be of essential scientific value. For example, satellite data from the 60's and 70's of the past century turned out to be essential for current research on global warming. Long-term preservation solutions of scientific and technical data is also needed, for example, in the automobile and aeroplane industry since construction data has to be retained for cases of insurance claims.

However, the driving force in development of solutions for long-term preservation of scientific data is to provide sustained global data access and exchange between globally distributed research collaborations [47–49].

III. OBJECTIVES

From the discussion of the previous sections we briefly formulate objectives for research and development in the field of long-term preservation of relational databases. A solution for the archiving of relational databases

- has to enable, to a good level of authenticity, permanent retention of the original internal data structure, the referential data integrity, and all technical and non-technical meta-information needed to keep the data technically accessible and intellectually understandable over the long term;
- must ensure that the data remains utilizable and processable by future data processing systems;

- must be able to completely detach databases from its proprietary database management software, hardware, and operating system environments;
- must completely rely on widely accepted and internationally standardized technologies;
- shall no longer require any specific software, maintenance, or administration for at least ten years, and reasonably longer if the full documentation of applied standard technologies remains available, provided that the physical bit-streams of all files remain intact;
- shall easily enable the reload to current and future relational database management system products, and thereby allow queries in almost the same complexity (on the database level) as in the original system from which the database was archived from;
- shall support the acquisition and standardization of non-technical data (which is usually not available from the physical database) across different business units and persons, and seamless integration with data and technical metadata.

IV. SOLUTIONS: SIARD WORKFLOW AND FEATURES

The method and Java application SIARD (“Software-Invariant Archiving of Relational Databases”) meets the aims outlined in Section III. The development of SIARD was a subproject of the strategic project ARELDA (Archiving of Electronic digital Data) of the Swiss Federal Archives and the Swiss Federal Administration.

SIARD consists of three stand-alone software applications, named A0, A1, and A2. Each application may be used by different people at different places during different stages of the workflow: database administrators and application responsables of the RDBMS, records managers of the data producers, IT professionals and general personnel of the archives, or archives customers. The output of each application is used as the input of the next stage’s application. Unfinished work can be saved any time and resumed later. All three applications are independent of the computer platform (cf. subsection IV D) and location. Communication between SIARD components and the RDBMS can take place either locally or through a TCP/IP network (which will be the usual case). All applications have XML-based, external configuration files that may also be edited manually.

Figure 1 summarizes the SIARD components and workflow. The RDBMS which contains the database to be archived is at the top of the figure. As examples, Figure 1 shows two commercial RDBMS products (Microsoft SQL-Server and Oracle) as well as an unspecific “other RDBMS” which could be an Open Source product, for example. (SIARD may also be used with Microsoft Access, that works best when Access is used in the SQL-Server compatibility mode.) Figure 1 sketches various set-ups for starting work with SIARD:

- Direct connection to the operational RDBMS.
- Migration of the database from the operational RDBMS to another RDBMS (including another instance of the same RDBMS product) using migration tools that are supplied by the vendors. (We denoted this intermediate RDBMS with “Oracle”, but it could be any other product.)
- Connection to the operational RDBMS through a vendor-supplied transparent gateway.

The first method (direct connection) is the usual case. However, direct connections of SIARD to the operational RDBMS may not be allowed or desired, for example due to security reasons or because the RDBMS is a high-availability system. (However, we emphasize that SIARD does *not* alter the RDBMS in any way! It solely performs read-only access operations on the RDBMS.)

But there are other possible reasons to add this intermediate step: For example, using transparent gateways or migration of the database from the original to another RDBMS may help to get more out of SIARD since it provides so-called “expert modes” for some specific products. These expert modes take into account and exploit product-specific features. Thus it may be advantageous to previously migrate the database from the original RDBMS to a RDBMS product for which SIARD provides an expert mode. Another reason could be that pre-processing of the database is required prior to archiving, for example data conversions or filtering. (SIARD can only exclude entire columns or tables from archiving.) Of course, such operations will never be performed on a production database.

As shown in Figure 1, the first SIARD application A0 analyzes and extracts a database from the RDBMS (with guidance by the user), and creates the archive files. After these tasks have been completed, the XML reference file (which contains all meta information about the results of the extraction) is loaded into application A1. In A1, the user adds further, mandatory and optional non-technical metadata on all levels of the object




FIG. 1: Overview of SIARD components and workflow: The relational database to be archived is extracted from a RDBMS (top), the SIARD archive is created for long-term preservation (middle) and finally restored in a RDBMS (bottom) for usage.

hierarchy (database, tables, columns etc.) as well as context metadata defined by the archival institution. Being complemented this way, the XML reference file is written back to the SIARD archive which is now ready for long-term preservation. The third application A2 comes into play when the archived database is requested for access and usage by a customer or required for other dissemination purposes. The reload is initiated through A2 by either the customer itself or by archive staff. We will describe the features of the three SIARD applications in more detail below.

Several tests have proven the applicability of the solution. The complexity of the tests performed so far range from tens of tables and a few thousand rows up to 250 tables and 250'000 rows from several commercial and open source database management products. The main difficulties were, as expected, proprietary, non-standard extensions to SQL in all RDBMS products, implementation-specific character sets, and extraction of some specific metadata from the system dictionaries, respectively.

A. A0: Analysis and Extraction

This is undoubtedly the core of SIARD since it determines the quality and extent of the database archive as far as technical aspects are concerned. A0 usually will be operated by the RDBMS application responsible (who has knowledge of the databases content), maybe assisted by a database system administrator (DBA) who has a deeper understanding of the technical background of the specific RDBMS. (Of course, A0 can also be used within the archival institution to migrate older non-SIARD data collections, for example.)

There are two ways to connect to the RDBMS with A0. The straight forward method is to allow A0 to connect as a DBA. Although A0 solely performs read-only operations on the database, this method is not recommended since it requires disclosure of the DBA password. The recommended method is that a DBA creates a new user in the RDBMS (named SIARD-A0, for example) and grants to it only those rights which are required by A0 to be operated properly. This also allows fine-tuning of privileges to an extent where A0 can only see exactly those parts of the database that are in fact subject to archiving².

At startup, A0 asks for the Java Database Connectivity (JDBC) [56] parameters to be used. The user either must enter the parameters manually into a panel, or else open up an XML-based A0 configuration file. (Usually, a DBA will provide a configuration file.) Afterwards, A0 asks for the access mode to be used. There are so-called “expert modes” for specific RDBMS products as well as a generic mode, used if no specific expert mode is provided by SIARD or the type of RDBMS is unknown. At the moment, there are only two expert modes (Oracle 7/8/9 and Microsoft SQL-Server 7/2000). Expert modes provide a broader range of database objects and metadata that can be archived since they exploit product-specific features. The generic mode has a rather narrow focus and solely uses the standard functionality of JDBC (which is continuously evolving from version to version, though.) New expert modes may be easily programmed and added to A0 without changes in existing code.

1. Automatic Analysis and Mapping

After successful connection to the RDBMS, A0 lists all³ schemata⁴ in the database and asks the user to select all or only some of them to be archived by SIARD. Afterwards, the database schemata (data dictionaries) are analyzed by A0's built-in SQL parser and validator. Additionally, A0 checks for data integrity and identifies, for example, isolated tables (which do not have any primary keys and no foreign or unique keys pointing to them).

During this process, A0 tries to automatically translate non-standard SQL constructs into forms that fully conform to the standard, provided that this conversion will not lead to any loss of information in the primary data. Otherwise the conversion is not performed. (The exact details of this criteria are described in the SIARD documentation.) Correspondingly, objects that do not conform to the SQL standard, nor can be automatically mapped to standard forms, are automatically set to the status “Cannot be archived”.

² Privilege tuning for A0 is described in the user manual and may depend on the specific RDBMS product used.

³ To be more precise: All schemata are listed only if A0 has logged in with DBA permissions. Otherwise, if A0 is logged in as an individual SIARD user, the list of schemata depends on the privileges that were granted by the DBA. Schemata and tables which are solely for internal purposes of the RDBMS are automatically ignored by the expert modes.

⁴ If the RDBMS supports catalogs, e.g. Microsoft SQL-Server, these are listed before the schemata.




FIG. 2: AO's workbench after database schemata (FLUGLE and SCOTT in this example) were analyzed. The right-hand pane shows the technical metadata and analysis results for the database object that is selected from the database object tree in the left-hand pane. Colored bullets and symbols indicate where user decisions are required to proceed.

When finished, the results are presented to the user as shown in Figure 2. There is a hierarchical, collapsible tree of database objects in the left-hand pane, while the right-hand pane shows the technical metadata from the data dictionary as well as the results of the analysis for the database object that is currently selected in the object tree. The tree of database objects comprises schemata, tables, table columns, one- and multiple-column primary key constraints, check constraints, triggers, views, the views' SQL code, view columns, stored procedure, functions, users, user roles, user role privileges, synonyms, and database links. There is either a small colored bullet or a document symbol attached to each object in the tree. Additionally, each root of a branch of the tree has assigned either a green check mark, a yellow triangle with an exclamation mark, or a red cross. The colors and symbols indicate the status of an object:

- A green bullet means "The object is or was made fully conforming with standard SQL and has proper data integrity – Ready for archiving".
- An orange bullet means "The object is or was made fully conforming with standard SQL but has problems with data integrity – Ready for archiving".
- A red bullet means "The object does not conform to and could not made conforming with standard SQL – Archiving is not possible without user intervention which may cause loss of information". Red bullets usually occur for unknown, proprietary or non-standard user-defined data types.

- A gray document symbol means “The object does not conform to and could not made conforming to standard SQL – A0 decided to exclude it from archiving.” This usually occurs when the SQL code of a view, constraint, trigger, stored routine, or function does not conform to standard SQL.
- A green check mark indicates that everything is okay (green bullets) or consolidated (gray document symbols) on subsequent nodes of the respective branch. This branch is ready for archiving.
- An exclamation mark indicates that there is at least one warning (i.e. a orange bullet) on subsequent nodes. The branch can be archived, though.
- A red cross indicates that there is at least one unresolvable problem (i.e. a red bullet) on subsequent nodes. A decision by the user is required (for example manual exclusion of the object from archiving).

If an object was automatically excluded or has a red bullet, the detected problems are explained in the “Details” tab of the right-hand pane.

2. Clearance

Archiving of the database is not possible as long as red crosses appear (i.e. any red bullets on lower levels of a branch), and the “Create Archive” button is disabled. Thus, the minimal action required by the user is to treat at least all objects with red bullets. To do so, the user has four possibilities. Three of them can be chosen either from the function button panel or the context menu of the right mouse button: First, the object may simply be manually excluded from archiving. In this case, the red bullet turns into a gray document symbol. However, the red bullet may come from an unknown data type which in fact is a valid standard SQL type but has a non-standard name (thus SIARD does not recognize the type). This usually should happen only in the generic access mode, but could also occur in expert modes when the version of A0’s expert mode is older than the version of the RDBMS. (In fact, many RDBMS use non-standard names for standard data types.) Pressing A0’s “Unknown and Proprietary Data Type” button, the user can define a catalog of synonyms for data type names to disclose to A0 the correct data type for this “unknown” data type name.

In the example shown in Figure 2, all red crosses are caused by a single problem which could not be resolved by A0’s initial automatic conversion algorithm: The data type “MY_TYPE” of the column “CODE” in table “PRINT_MEDIA” is an unknown data type and thus could not be automatically converted to a standard SQL data type without loss of information. In consequence, A0 assigned it a red bullet, and the analysis result (visible in the right-hand pane) says “SQL3 Check: Type conversion to SQL3 impossible”. A0 did not assign a gray document, however, since there may be still the chance to solve a possible naming conflict by a user-supplied synonym rule as described above. In this example, the data type “MY_TYPE” is a user-defined data type with a non-standard definition. However, inspection of the original database would reveal that (in this example) “MY_TYPE” can be mapped onto the standard SQL data type “varchar(10)” without loss of information. The user can easily resolve the problem by defining a naming rule “MY_TYPE → VARCHAR(10)” using A0’s “Unknown and Proprietary Data Type” button.

The second possibility is that the red bullet in fact comes from a proprietary non-standard data type. The user may either exclude the object from archiving (third possibility) or use the “Proprietary Data Type” button to define a deep conversion to a standard data type. This may cause loss of information, though, and is only possible in expert modes⁵ (since the input and output of a conversion must be well defined). The fourth possibility is to perform a conversion of unknown or proprietary data types within the original database (i.e. outside of SIARD’s A0 application), for example by using vendor-supplied specific SQL CAST functions⁶.

All conversions initiated in A0 are *never* performed in the original database but only after the extraction of the database from the RDBMS. We also emphasize that exclusion of an object from archiving does not mean that the object will be invisible in the final archive. In contrast, all meta information about the object (available from the original database) will be documented in the archive (including its non-standard SQL code). But it means that the object (for example a table, table column, check constraint, stored routine, or trigger) cannot be actively restored anymore in a RDBMS later.

⁵ This conversion functionality is not yet implemented in the current release of SIARD.

⁶ A similar procedure may be applied to non-standard SQL code in views, triggers etc. However, we think that it is not desirable to change a database too deeply for the purpose of archiving. Furthermore, many RDBMS use a proprietary procedural programming language for stored procedures which cannot be easily mapped onto the SQL/PSM language (cf. subsection II A 1).

Of course, the user may also exclude any valid object (with a green bullet) from archiving, for example if archiving is not desired due to appraisal decisions. The excluded object will still be fully documented in the SIARD archive (though its data content will be excluded). We emphasize that automatic and manual exclusion of objects (for example entire tables or single table columns) is only possible because A0 restores a proper data integrity after any exclusion. For example, when a table is excluded from archiving, the user is warned if there are any primary keys in this table, or if any foreign or unique keys of other tables are pointing to that table. If the user confirms exclusion, A0 removes the key constraints from the referencing tables (without a cascading deletion of the table rows) and automatically also excludes from archiving all views, triggers, check constraints etc. which contain references to the excluded table. A similar procedure applies if single table columns are excluded. Finally, we note that any exclusion operation can be reversed again by using A0's 'Undo' function.

Orange bullets primarily indicate isolated tables (with no key constraints in it and no foreign keys pointing to it) and do not require user intervention. However, using the context menu of the right-mouse button or the function pane, the user may define its own primary as well as foreign and unique key constraints manually⁷ on any table, assisted by an interactive panel. A similar function and panel is provided to define user-added check constraints (using standard SQL code only). Both possibilities may be optionally used if linkage information or check constraints for the table do implicitly exist but are hidden in external application software which operates on the original database. Adding the corresponding information from the external application's system documentation will complement and improve the database archive while external software usually will not be archived. (Note that user-added key and check constraints are only defined within the SIARD archive, while the original database is not altered.)

During work, the user can save intermediate states at any time, and resume his work later. Furthermore, the application A0 has a special added-on object at the end of the object tree, named 'changelog'. In fact, A0 traces and remembers all changes to the original state of the database. This includes activities performed automatically by A0 during its initial analysis of the original database, as well as all changes caused by manual operations of the user later on. Each entry in the 'changelog' contains a time stamp, a short description of the activity, and the nature of the change. The log file will be part of the SIARD archive (and be extended further in the subsequent module A1 of SIARD). This enables re-tracing of the archiving process at any time in the future and thus supports authenticity of the database archive.

3. Creation of the Archive

The "Create Archive" button is enabled when the database is ready for archiving. This happens as soon as all root nodes of branches in the object tree have assigned either green checkmarks or yellow exclamation marks. The latter indicate warnings (orange bullets) on subsequent nodes (for example isolated tables) which may be acceptable. For the example shown in Figure 2 this state can be reached by either defining the before-mentioned synonym rule for the data type "MY_TPE", or by exclusion of the table column "CODE" from archiving (while its metadata will still be included in the SIARD archive).

When pushing the "Create Archive" button, A0 asks for the location where the archive shall be saved. Afterwards, A0 starts to load the primary table data from the original database, performs all necessary conversion operations on them⁸, and creates all SIARD archive files as depicted in Figure 1:

- **SQL-DDL files:** These files contain standard SQL:1999 Data Definition Language (DDL) statements only. Together they represent the definition of a self-contained relational database, comprising all objects and attributes of the original database, except for those that were excluded.
- **Table data files:** These are files which contain the primary data (except for large objects, see below) of the database defined in the DDL files above. There is one file per table. The data of one table row is contained in one line of the file, and data items have variable lengths. Rather than putting absolute delimiters between two adjacent data items, we use a simple algorithmic token for delimitation⁹.

⁷ In contrast to objects which belong to the original database, user-added key and check constraints cannot be excluded from archiving but only deleted completely (using the 'Delete' function or button). The delete function is disabled for all other objects.

⁸ Depending on the amount of data and the network transfer capacity, this may require a few seconds up to many hours. Primary data is loaded and processed sequentially and will not necessarily have to be stored on A0's local machine. In fact, the RDBMS, the application SIARD A0, and the SIARD archive files may all be at different remote locations and connected by a network.

⁹ Every item has the form "*l*, ...;", where *l* is the number of characters of the data item "...". (The semicolon is not required but

- **Large Object Files:** Each file contains a single data item of a so-called large object string type (provided that there really exist such data types in the original database, of course). This is either a character large object¹⁰ (CLOB) or a binary large object (BLOB) which was embedded in a table row of the original database. The former type is an arbitrary sequence of characters (for example a narrative document encoded in XML), the latter is an arbitrary sequence of bytes (for example an image file). Both may be up to several gigabytes in size. BLOB files just contain a hexadecimal dump of the original BLOB, thus may contain anything¹¹.
- **XML Reference File** This XML document contains all information from A0's workbench. More precisely, it contains three kinds of metadata: Firstly, the complete database logic that is also contained in the SQL-DDL files (but encoded in XML). Secondly, all metadata from those database objects that were excluded from archiving (and thus do not appear in the DDL files), including the code of stored routines, triggers, views etc. Thirdly, the data from the "changelog" which reveals when and what actions were performed during the archiving process, either automatically by A0 itself or manually by the user.

Note that a SIARD archive is made up of all files listed above, not only of the DDL files (which only contains those parts of the original database which may be reloaded into another RDBMS again). In particular, the XML Reference file accomplishes the smooth integration of archived database objects, excluded objects, and all kinds of metadata. All files are plain text files (but may contain hexadecimal strings for binary data), and Unicode/UTF-16 encoding [25–27] is used to overcome implementation-defined character encodings of the original RDBMS, and to preserve multilingual character sets (including non-latin alphabets).

The XML reference file may also be exploited for various other tasks. For example, it may be used by XML schema mapping tools to generate metadata subsets for import into finding aid or catalog systems. Or it may be used in the future for easy migration of the SIARD archive to other formats, for example to a future release of SQL (though SQL is developed upward compatible).

Furthermore, for more convenient handling by humans only, every table data file has a short XML header that contains the main table and column metadata such as names, data types, sizes, key constraints (primary, foreign, and unique), or default values. These table data files will not be altered anymore during the next steps of the SIARD workflow and therefore may be right away utilized for other purposes. For example, they can be viewed with any trivial text editor program. In addition, every SIARD archive contains a XML stylesheet language (XSL) file named "dmpFile.xsl" in its data file directory. Therefore, if a table data file is opened with a web browser¹² from within this directory, the user automatically gets a pretty-print version of the data file (rather than clumsy to read raw data rows), including named columns as well as vertical and horizontal table lines. For example, the user may print this version (or catch the HTML output) for the purpose of non-technical distribution of the data.

The SIARD archive has a fixed structure of file directories. This directory structure as well as some other information is contained in an XML file named "archiveInfo.xml" that is located in the root directory of each archive (and may be used by other software programs).

B. A1: Description

After creation of the SIARD database archive, the next step shown in Figure 1 is to add complementary metadata which is not available from the original database or the RDBMS. As discussed in Section II B,

used for convenient reading by humans.) This representation is foolproof (whereas absolute delimiters are not) and independent of the character encoding (whereas the byte length of the data item strongly depends on it, in particular for variable-width multi-byte encodings like, for example, UTF-8). This approach requires that every line is processed sequentially from the start, which is not a disadvantage.

¹⁰ Actually, before writing the CLOB files, SIARD A0 converts the CLOBs of the original database into National Character Large Objects (NCLOB) via translation of all characters (which are stored in the implementation-defined character encoding of the RDBMS) to the fixed-length Unicode character encoding UTF-16. Otherwise, the characters in CLOB files would be rather useless without exact knowledge of the original, implementation-dependent character encoding.

¹¹ The "expert modes" of SIARD A0 also use CLOBs and BLOBs to accommodate certain kinds of proprietarily typed large objects without loss of information. For example, the expert mode for Oracle puts an Oracle BFILE (i.e. a BLOB that is stored outside of the RDBMS) and Oracle LONG RAW objects into BLOBs, while the expert mode for Microsoft SQL-Server puts Microsoft TEXT and IMAGE objects into NCLOBs and BLOBs, respectively.

¹² This requires a browser which supports XSL, for example current releases of Mozilla [53] or Microsoft Internet Explorer




FIG. 3: AI’s workbench to add complementary narrative and context metadata to the database archive. For a single database object (selected in the object tree on the left-hand side) metadata is entered in the right-hand pane, while context metadata on the database level is entered in a multi-tab “Context Metadata Editor” window (shown in front of the right-hand pane).

this kind of metadata is indispensable to enable long-term intelligibility and comprehensibility of the archived database.

Such metadata is added to the database archive using application SIARD A1. The users of A1 will not have to be the same as those of A0. Instead, it may be more non-technical staff as, for example, records managers, data asset managers, experimental scientists, or application responsables. A1 only reads and modifies the XML reference file created by the application A0 (cf. Section IVB), and the state of work may be saved at any time and resumed later. Thus the complementary metadata for A1 may be gathered across different business units by simply forwarding this XML file from one person to the other.

Figure 3 shows the graphical user interface of A1. Again, there is a database object tree on the left-hand pane. It is basically the same as in A0, except for the colors which now have a different meaning: red crosses indicate that there are mandatory but not yet filled in metadata fields on subsequent nodes, while green checkmarks indicate that all mandatory metadata has been provided on subsequent nodes.

For a database object that is selected in the object tree, the right-hand pane shows the metadata that was contributed by A0 (and thus comes from the original database), plus metadata that was added using A1. On the object level, the user can enter metadata at three locations in the right-hand pane: An arbitrary full text object name which may be more meaningful than its technical name (which is often an abbreviation), an arbitrary narrative description of the object which may improve the intelligibility of the object, and finally a user-added code list. The latter may be essential since code lists are frequently documented only outside a database. In our example in Figure 3, the column `GRADE` of the table `STUDENT.SCHEDULE` contains grades that were assigned to students after visiting certain college courses[52]. The grades are coded with one or two characters (e.g. A+ or F-). However, the original database does not contain any information about what these codes

mean. Therefore, the user has added a code list (using the “Add/Edit Value List” button) which explains that, for example, A- means “85 - 90% of the exam questions were answered correctly by the student”. Often codes are even less self-explanatory (imagine a code “92” which in fact means “application rejected”).

The context metadata on the database level is added and edited in a separate, multi-tab “Context Metadata Editor” panel (see Figure 3). It contains various tabs, each of them covering a certain subject in the context of the database’s usage, creation, original IT environment, history of the system, data authority, or provenance. Each tab contains a number of metadata fields with colored buttons that indicate whether filling in is mandatory or optional. A field is either a free text field or a pull-down list to select from predefined values. If the user points into a field, a description of the field is shown at the top of the panel. In addition, the descriptions can be displayed in different languages (which can be selected at the upper-right corner of the panel).

The extend and structure of the context metadata usually depends on specific requirements of the archival institution. Therefore, the “Context Metadata Editor” (CME) panel is fully customizable by using SIARD’s “Context Metadata Schema Editor” (CMSE) which is part of A1. Actually, the CME is built from an XML file which contains the CME layout, and the CMSE is basically a graphical user interface to manipulate this XML file. The user may define its own CME panel, including individual tabs as well as individual metadata fields and descriptions (using arbitrary languages). Thus it is possible to define several standardized context metadata schemata (CME panels) for different classes of databases or archival scopes. Depending on the class or purpose, the appropriate CME to be loaded in A1 to add customized context metadata to the SIARD archive.

In addition, there is a special CME tab “System Documentation” which allows the user to enclose arbitrary files to the context metadata part of a SIARD archive¹³. These may be, for example, PDF documents or TIFF images that are taken from the original RDBMS and database documentation, for example system and user manuals, log files, security reports, original data dictionaries etc. (One could even think of MPEG video documents showing people at work with the original production database system.)

While the user is working with A1, the application records all user modifications in the subnode “A1” of the “changelog” which was already mentioned in Section IV A 3. Furthermore, for all context metadata entered through the CME panel, A1 does not only store the filled in data values of a metadata field but also the entire, language-specific version of the field description which was visible in the panel at that moment when the user entered or modified the data value. Making the field descriptions and the mandatory/optional characteristic integral parts of the SIARD archive is necessary since metadata catalogs (i.e. CME panel configurations) may change or evolve over time, and the meaning of single metadata items may change.

At any time, the user may convert the current version of the XML reference file into an HTML document for convenient review or distribution. When all mandatory metadata has been filled in, the “Finalize” button is enabled. When this button is pushed, A1 will warn the user that no further changes to the reference file will be possible later, and will then write the final version of the XML reference file. This version now contains all metadata that was generated or collected by A0 and A1. The SIARD database archive has been completed and is now suitable long-term preservation. It does not depend on any specific hardware and software (not on SIARD as well), it consists of plain text files only (except for optional, user-added PDF or TIFF files), and it is solely based on technologies which are widely used and internationally standardized.

C. A2: Reload and Access

Apart from their primary purpose, SIARD database archives may also be utilized and processed by other software applications, for example in data dissemination and exchange or data warehouses. Because of its open technologies and its high level of standardization, SIARD archives may also be easily converted into forms that can be accessed directly in the World Wide Web. However, this may require some additional, expensive hardware and complex software (e.g. an application web server).

We have therefore added a third application to SIARD, named A2, which does not require additional infrastructure at all. The same infrastructure that was used for archiving the database (i.e. creating a SIARD archive) will be perfect. Basically, A2 is a simple RDBMS client but enables users to reload SIARD archives into a RDBMS “on demand”, and then provides multi-user remote network access for querying and browsing the data together with its technical and descriptive metadata in one graphical user interface.

¹³ Note that A1 does not inspect these files at all. It’s up to the user’s responsibility to use document and data formats that are suited for long-term preservation.




FIG. 4: A2's Data Browser after reload of two SIARD database archives into an Oracle database. The pane at the bottom shows the data of the table which is selected in the database object tree (left-hand pane), while the pane at the top shows its metadata. The other database ("BILLY") is accessed by selecting the second tab of the left-hand pane.

We emphasize that A2, in its current version, is only a prototype (thus still has some bugs, limited functionality, and some security deficits), and currently only works with an Oracle RDBMS as the reload target system. However, it will be rather simple to adapt A2 to work with other RDBMS products as well.

Working with A2 requires a running standard Oracle database instance (reachable either through a TCP/IP network or the local host) and some very simple preparations prior to operate A2. This preparation, however, has to be done by a system administrator: A new tablespace to hold the restored SIARD schemata and data should be created (and provide enough space for the estimated SIARD restore operations). In addition, a user with specific rights (which are described in the user manual) has to be added to the RDBMS. This user will act as a "SIARD archive manager" to control and serve connections from A2 clients. Its database schema will be the repository for information about all reloaded archives, registered users that are allowed to reload SIARD archives to this RDBMS, and it records reload operations and connection times of individual users. This schema will be initialized automatically as soon as any SIARD A2 client is connecting to this RDBMS for the first time.

A2 is controlled by an XML configuration file which contains the profiles of one or more RDBMS that are available as targets for SIARD restores. A profile identifies the Oracle database instance, the tablespace used for restore operations, and the before-mentioned "SIARD archive manager" user. Moreover, it provides additional information required by A2 to connect to the RDBMS. This configuration file will be pre-configured and provided to end-users of A2 by the RDBMS system administrator.

After starting A2, the user is asked to choose (from a file selector panel) the "archiveInfo.xml" file from the SIARD database archive to be reloaded, to select one of the pre-configured connection profiles, and to decide whether or not large object string files (BLOB and NCLOB, cf. Section IV A 3) shall be reloaded too. Finally, the user either pushes the "Create New User" button to create a personal database account for working with

A2, or else chooses an existing account from a pull-down list. After entering the user password¹⁴, A2 reloads all schemata contained in the SIARD database archive¹⁵. The user may sequentially reload as many SIARD database archives as needed and use them all together in a single A2 session. If a SIARD database archive was already reloaded by another user, it will not be reloaded a second time. (Users of A2 only have read access to the reloaded database.)

Figure 4 shows the main component of A2, the “Data Browser” panel. Again, there is the database object tree navigator already known from A0 and A1. In addition, there are tabs at the top of the tree panel to switch between different reloaded databases. The pane at the bottom shows the data of the table which is selected in the database object tree, while the pane at the top shows its metadata. Further metadata is found on subsequent nodes and in the “Context Metadata” branches below the “archives” root node. For example, the code list for student grades (which was user-defined during work with A1 in Section IV B) can be accessed on the corresponding leaf node “code list” in the FLUGLE schema (table STUDENT_SCHEDULE, column GRADE).

As can be seen in Figure 4, there are some (but not all) views reloaded. These views (marked with a green bullet) are actually proper views, and selecting them will show the view’s data in the data pane. The other views (those with a grey document symbol) do not show any data because they were excluded from archiving by SIARD’s application A0 due to non-standard SQL code in the original database. However, the view’s definition as well as the reason why it was excluded will be visible in the upper metadata pane.

The user may browse the database, select parts of table data in the data pane, and export it as comma separated (CSV) files. Furthermore, users who are familiar with the SQL query language may use the “SQL” tab (at the top of the upper pane) to switch to a separate panel where arbitrary SQL:1999 query statements can be composed and send to the Oracle RDBMS. The results of the query will be shown in a second pane where they can be selected all or in parts for export to a CSV text file.

If a user A decides to quit A2, he or she will be asked if the restored database shall be deleted. However, if there is still another user B currently registered for using the same database, it will not be deleted (but the registration of user A for this database will be removed).

In conclusion, SIARD A2 seamlessly integrates all metadata and data (whereas the reloaded database only contains the metadata from the DDL files described in Section IV A 3), and it enables the user to perform simple as well as complex SQL queries on restored databases. Query results can be exported to the local machine. Several A2 clients may connect simultaneously to the same restored database, and A2 provides controlled multi-user remote access¹⁶ to SIARD database archives.

D. Development Environment

Software development of SIARD was carried out by the Swiss Federal Archives together with Trivadis (Switzerland) AG [54]. The SQL parser and validator as well as the context metadata schema editor was developed by one of the authors (SH).

The SIARD software is platform independent, relying on the programming language Java and the Java virtual machine as an interface to the operating system. It was tested under Solaris 7 and 8, Red Hat Linux 7, and Windows NT / 2000 / XP. The JDBC driver may need a specific environment to run properly (e.g. an MS Windows for a Windows specific authentication on an MS SQL-Server), but these restrictions are solemnly defined by the driver.

Currently, Eclipse [55] is used as Integrated Development Environment, mainly because of its file-based approach. The technologies used in SIARD are

- Java 2 Software Development Kit (J2SDK) 1.4 by Sun Microsystems [56].
- Java Foundation Classes / Swing [56] were used for the graphical user interface.

¹⁴ Warning: This user password only authenticates the human user of A2, not the connection between A2 and the RDBMS itself, whereas A2 connects to the RDBMS as the “archive manager” user which has granted DBA rights, and its password is contained as clear text in the A2 configuration file! This is a severe security risk. As mentioned: A2 is still a prototype, and a production release of A2 will have to contain a secure authentication of A2 clients. Thus we *strongly* recommend *not* to use A2 with a production database (i.e. one which holds any important data aside from SIARD database reloads).

¹⁵ Depending on the size of the archive and the network transfer capacity, this may require a few seconds up to many hours.

¹⁶ Any Firewalls between the A2 client and the RDBMS will have to be properly configured, though.

- Java Database connectivity (JDBC) 3.0 [56] for data retrieval from and reload to the database management systems.
- Java API for XML Processing (JAXP) 1.2 [56] which handles operations on XML data (not needed anymore for J2SDK 1.4.2 and higher).
- The Extensible Markup Language (XML) 1.0 [22] for semantic markup of data in external text files created and read by SIARD components.
- The Extensible Stylesheet Language Family (XSL): XSL Transformations (XSLT) Version 1.0 [58] is a XML-based declaration language for displaying and transforming XML files.
- XML Schema (XSchema) 1.0 [24] provides automatic consistency and integrity checks on the SIARD XML files.
- The Structured Query Language (SQL) ISO/IEC 9075:1999 [59] for definition of database layouts.
- The Unicode Transformation Format UTF-16-UCS-2 [25–27] for platform independent and multilingual character encoding in all SIARD text files.
- Oracle 8/9i and Microsoft SQL-Server 7/2000 relational database management systems and Microsoft Access 97/2000 were used for testing SIARD.

The JDBC driver implementations of the different database manufacturers provide a varying degree of compliance with the JDBC specifications. Especially, the functions for querying the database for metadata leave much to be desired. This metadata could be extracted from the database by database specific SQL-like queries, which is why the database access in SIARD is encapsulated in so-called “modes”, allowing for database-dependent enhancements.

“Expert modes” allow manufacturer-specific access to the database engines. Today, expert modes have been implemented and tested for JDBC Drivers for the Oracle and Microsoft products mentioned above. For the use with other RDBMS products, a generic mode is provided. The open architecture of SIARD allows the simple development of further expert modes for archiving from other database products since they can be added dynamically (i.e. without changes in existing code).

As explained in Section IV A 3, SIARD ultimately produces plain text files only. Testing conformance of the SIARD XML files with the XML standard is simple: all have dedicated XML schemata and thus can be validated, and there exist a multitude of XML parsers which provide just this functionality. The compliance of the SQL structures files with standard SQL:1999 is governed by SIARD’s own, plugged-in SQL parser and validator that checks the syntax of the SQL expressions as well as the dependencies between these expressions. (For example, for a table to be created in a specific database schema, the schema must have been created first.) Independent cross-checking of SIARD’s SQL files is possible too, though we are aware of only one other tool which provides broad SQL validation functionality [34, 50].

V. CONCLUSIONS

We have discussed problems and relevance of long-term preservation of relational databases for usual archival institutions like national archives and scientific data archives that have to ingest data from a broad diversity of vendor products. We have argued that the common current ingestion and preservation practices may suffice for ingestion of small and simply structured data sets but suffer from insufficient integration of data and metadata, lack of automation in the ingestion of large amounts of data, error-proneness, and in general require extensive manual effort and intervention to make the data accessible and usable. Without having more efficient solutions at hand, the rapidly growing size and complexity of relational databases in modern relational database management systems will rapidly outpace the ability of archives to ingest, manage, and preserve them.

Furthermore, we have surveyed the relation between present-day database management system products and the standardized data definition, query and manipulation language SQL, and have critically appreciated its applicability to long-term preservation of databases. From this discussion we concluded that widespread non-standard, vendor-supplied enhancements and additions do not allow for one-to-one ingestion from such systems, and that no “SQL for Archiving” exists. Nevertheless, standard SQL may be reasonably exploited for long-term preservation purposes when data and data logic are actively extracted from database management systems by specialized ingest tools which map different “SQL flavors” to generic SQL, and transparently trace and document those parts which cannot be mapped.

Finally, we have presented the method and platform independent application “Software-Invariant Archiving of Relational Databases” (SIARD), developed at the Swiss Federal Archives. It is an efficient, traceable and controllable ingest tool to detach relational data from any specific hardware and software environment, and thereby enables, up to a reasonable level, retention of its original authenticity, integrity, accessibility, and usability. SIARD integrates data with data logic and descriptive metadata and supports the intelligibility of databases for long-term archival preservation and access.

Readers who are interested in testing and reviewing SIARD or writing new expert modes may contact the authors. SIARD is property of the Swiss Federal Administration.

-
- [1] J. Garrett, J. and D. Waters. Preserving Digital Information: Report of the Task Force on Archiving of Digital Information, commissioned by the U.S. Commission on Preservation and Access and The Research Libraries Group, Inc. (1996), <ftp://ftp.rlg.org/pub/archtf/final-report.pdf> or <http://www.rlg.org/ArchTF/tfadi.index.htm>
 - [2] Preserving Scientific Data on Our Physical Universe: A New Strategy for Archiving the Nation's Scientific Information Resources. U.S. National Research Council, National Academy Press, Washington, DC (1995), ISBN 030905186X, <http://www.nap.edu/catalog/4871.html>
 - [3] K. Thibodeau, Preserving Scientific Information on the Physical Universe. IASSIST Quart. 19(4) (1995) pp. 25-29, <http://datalib.library.ualberta.ca/iassist/publications/iq/iq19/iqv194thibodeau.pdf>
 - [4] J. Rothenberg, Ensuring the longevity of digital documents. Scientific American, **272** (1), 1995, 42-47.
 - [5] J. Rothenberg, Avoiding technological quicksand: finding a viable technical foundation for digital preservation – A Report to the Council on Library and Information Resources. Washington, D.C.: Council on Library and Information Resources, 1999. (<http://www.clir.org/pubs/reports/rothenberg/contents.html>)
 - [6] The Long-term Preservation of Authentic Electronic Records: Findings of the InterPARES Project – Task Force Reports, The International Research on Permanent Authentic Records in Electronic Systems (InterPARES) Project 1 (1999 - 2001), <http://www.interpares.org/book/index.cfm>
 - [7] For an overview of literature dated before 2000, the interested reader is referred to "Preservation of electronic information: a bibliography" by Michael Day, <http://homes.ukoln.ac.uk/~lismd/preservation.html>
 - [8] Report of The Digital Archive Directions Workshop on June 22-26, 1998, hosted by the U.S. National Archives and Records Administration, College Park, MD (<http://ssdoo.gsfc.nasa.gov/nost/isoas/dads/ws.html>)
 - [9] US and International Efforts Towards ISO Archiving Standards, <http://ssdoo.gsfc.nasa.gov/nost/isoas/overview.html>
 - [10] UNESCO Charter on the Preservation of the Digital Heritage, adopted at the 32nd session of the General Conference of UNESCO, 17 October 2003, http://portal.unesco.org/ci/ev.php?URL_ID=13366&URL_DO=DO_TOPIC&URL_SECTION=201&reload=1070274443
 - [11] Final Report of the U.S. National Science Foundation Workshop on Research Challenges in Digital Archiving: Towards a National Infrastructure for Long-Term Preservation of Digital Information, April 12-13, 2002, Warrenton, Virginia (August 2003), <http://www.si.umich.edu/digarch/>
 - [12] The State of Digital Preservation: An International Perspective, Conference Proceedings, Documentation Abstracts, Inc., and Institutes for Information Science, Washington, D.C., April 24-25, 2002, Council on Library and Information Resources, Washington, D.C. (2002), ISBN 1-887334-92-0, <http://www.clir.org/pubs/reports/pub107/pub107.pdf>
 - [13] For an overview, see: The Cedars Guide to Digital Preservation Strategies. CURL Exemplars in Digital ARchives (Cedars) Project, <http://www.leeds.ac.uk/cedars/guideto/dpstrategies/dpstrategies.html>, and the Electronic Resource Preservation and Access Network (ERPANET) Project of the European Union, <http://www.erpanet.org>
 - [14] R. F. Sproull and J. Eisenberg (editors), Building an Electronic Records Archive at the National Archives and Records Administration – Recommendations for Initial Development. U.S. National Research Council, The National Academies Press, Washington, D.C. (2003), ISBN 0309089476, <http://books.nap.edu/catalog/10707.html>
 - [15] ISO 14721:2003 – Space data and information transfer systems – Open archival information system – Reference model. International Organization for Standardization, Geneva, Switzerland (2003). Except for cover pages and appendices, this standard is identical with: CCSDS 650.0-B-1 Recommendation for Space Data System Standards, Consultative Committee for Space Data Systems, National Aeronautics and Space Administration, Washington, D.C. (2002), <http://ssdoo.gsfc.nasa.gov/nost/wwwclassic/documents/pdf/CCSDS-650.0-B-1.pdf>.
 - [16] Report of the ERPANET workshop "The Long-term Preservation of Databases", April 9-11, 2003, Bern, Switzerland, The Electronic Resource Preservation and Access Network Project (2003), <http://www.erpanet.org/www/products/bern/bern.htm>
 - [17] Readers who are unfamiliar with relational databases and SQL may find a competent introduction in both subjects here: J. Melton and A. R. Simon, SQL:1999 – Understanding Relational Language Concepts. Morgan Kaufmann Publishers, Imprint of Academic Press, San Diego (2002). Also see E. F. Codd's famous article in Ref. [18].
 - [18] E. F. Codd: A Relational Model of Data for Large Shared Data Banks, Communications of the ACM, Vol. 13, No.

- 6 (1970), pp. 377-387, <http://www.acm.org/classics/nov95/toc.html>.
- [19] ERPANET Case study Pharmaceuticals, Electronic Resource Preservation and Access Network – ERPANET (2003), ISSN 1741-8682, <http://www.erpanet.org>
 - [20] ISO/IEC 9075:2003 – Information technology – Database languages – SQL. International Organization for Standardization, and International Electrotechnical Commission. Geneva, Switzerland (2003).
 - [21] A. Eisenberg et al., SQL:2003 Has Been Published. ACM SIGMOD Record, Vol. 33, No. 1 (2004), <http://www.acm.org/sigmod/record/issues/0403/E.JimAndrew-standard.pdf>
 - [22] Extensible Markup Language (XML) 1.0 (3rd edition) and 1.1, World Wide Web Consortium Recommendation (2004), <http://www.w3.org/XML/>
 - [23] A. Eisenberg and J. Melton. SQL/XML is Making Good Progress. ACM SIGMOD Record, Vol. 31, No. 2 (2002), <http://www.acm.org/sigmod/record/issues/0206/standard.pdf>
 - [24] XML Schema 1.0, World Wide Web Consortium Recommendation (2001), <http://www.w3.org/XML/Schema>
 - [25] The Unicode Standard, The Unicode Consortium, <http://www.unicode.org/standard/standard.html>. The Unicode Standard is aligned with the Universal Multiple-Octet Coded Character Set (UCS), ISO/IEC 10646, see Ref. [26].
 - [26] ISO/IEC 10646:2003 Information technology – Universal Multiple-Octet Coded Character Set (UCS). International Organization for Standardization, and International Electrotechnical Commission, Geneva, Switzerland (2003).
 - [27] RFC 3629 (2003) – UTF-8, a transformation format of ISO 10646, <ftp://ftp.rfc-editor.org/in-notes/rfc3629.txt>, and RFC 2781 (2001) – UTF-16, an encoding of ISO 10646, <ftp://ftp.rfc-editor.org/in-notes/rfc2781.txt>
 - [28] ISO/IEC 13249 – Information technology – Database languages – SQL Multimedia and Application Packages. International Organization for Standardization, and International Electrotechnical Commission. Geneva, Switzerland (2003)
 - [29] Differences in the implementation of SQL by various products are well described in (for example): K. E. Kline with D. Kline and B. Hunt: SQL in a Nutshell, 2nd Edition, O'Reilly, Sebastopol (2004), ISBN 0-596-00481-8; J. S. Bowman, S. L. Emerson, M. Darnovsky: The Practical SQL Handbook: Using SQL Variants, 4th Edition, Addison Wesley Professional (2001), ISBN 0201703092; M. Schmidt and T. Demmig: SQL Ge-packt, MITP, Bonn (2001), ISBN 3826606868; P. Gulutzan: Standard SQL – Do IBM, Microsoft, and Oracle support the SQL:1999 standard?, DBazine.com, <http://www.dbazine.com/gulutzan3.shtml>. Free resources are, for example: <http://troels.arvin.dk/db/rdbms/>, <http://dev.mysql.com/tech-resources/features.html>
 - [30] PDF-Archive Committee, The Enterprise Content Management Association (AIIM), http://www.aiim.org/pdf_a/; also see: L. Rosenthal: PDF/A – PDF for Archiving. Presentation at the PDF Conference 2004, Orlando, USA, http://www.pdfsages.com/downloads/PDF_for_Archive.pdf.
 - [31] PDF Reference 3rd edition – Adobe Portable Document Format Version 1.4: Adobe Systems Incorporated, published by Addison-Wesley (2001), ISBN 0-201-75839-3, http://partners.adobe.com/asn/acrobat/docs/File_Format_Specifications/PDFReference.pdf. Other versions at: <http://partners.adobe.com/asn/tech/pdf/specifications.jsp>
 - [32] ISO 15930 – Graphic technology – Prepress digital data exchange – Use of PDF. (PDF/X) International Organization for Standardization, Geneva, Switzerland (2001-2003).
 - [33] Federal Information Processing Standards Publication (FIPS PUB) 127-2: Database Language SQL, U.S. National Institute of Standards and Technology (1993), <http://www.itl.nist.gov/fipspubs/fip127-2.htm>. Download the FIPS 127 test suite at http://www.itl.nist.gov/div897/ctg/sql_form.htm. Also see: FIPS 193: SQL Environments, U.S. National Institute of Standards and Technology (1995), <http://www.itl.nist.gov/fipspubs/fip193.htm>.
 - [34] “Mimer SQL-Validator” by Mimer Information Technology (<http://www.mimer.com/>). This commercial C++ program is used to provide a free web service for SQL validation covering SQL-92, SQL:1999, and SQL:2003: <http://developer.mimer.com/validator/>. Also see Ref. [50]. We have not yet tested this product. SIARD has its own, built-in SQL parser and validator.
 - [35] Research Libraries Group: Trusted Digital Repositories – Attributes and Responsibilities. An RLG-OCLC Report. RLG Mountain View, California, USA, Research Libraries Group (2002), <http://www.rlg.org/longterm/repositories.pdf>
 - [36] D. O. Stephens and R. C. Wallace, Electronic Records Retention: New Strategies for Data Life Cycle Management, ARMA International, Lenexa, USA (2003), ISBN 1-931786-08-9
 - [37] Ch. M. Dollar, Authentic Electronic Records: Strategies for Long-Term Access, Cohasset Associates Inc., Chicago, USA (2002), ISBN 0-9700640-0-4
 - [38] T. Zürcher Thrier: Sicherung und Bewertung von Datenbanken im Schweizerischen Bundesarchiv. Arbido No. 3/2003, pp. 19 - 21, VSA/BBS/SVD (2003)
 - [39] See, for example: Public Record Office Victoria: Victorian Electronic Records Strategy (VERS) – Advice to Victorian Agencies, Advice 7: Preserving Records in Databases. Public Record Office Victoria, Australia (2003), <http://www.prov.vic.gov.au/vers/standards/PROVRMAdvice7.pdf>.
 - [40] E. Shepherd and Ch. Smith: The application of ISAD(G) to the description of archival datasets. Journal of the Society of Archivists, Vol. 21, No. 1 (2000), pp. 55 - 86. Also see: E. Shepherd: The Application of ISAD(G) to the Description of Archival Datasets, Presentation at the ERPANET Workshop on Long Term Preservation of Databases (cf. Ref. [16]), [http://www.erpanet.org/www/products/bern/presentations/WS%20Bern%](http://www.erpanet.org/www/products/bern/presentations/WS%20Bern%20)

- [41] See, for example: Assessment of the Usefulness and Availability of NASA's Earth and Space Science Mission Data. U.S. Research Council, National Academy Press, Washington, D.C. (2002), ISBN NI000398, <http://www.nap.edu/catalog/10363.html>
- [42] See, for example: Study on the Long-term Retention of Selected Scientific and Technical Records of the Federal Government: Working Papers. U.S. National Research Council, National Academy Press, Washington, D.C. (1995), ISBN NI000157, <http://www.nap.edu/catalog/9478.html>
- [43] J. M. Esanu and P. F. Uhler (editors), The Role of Scientific and Technical Data and Information in the Public Domain. U.S. National Research Council of the National Academies, The National Academies Press, Washington, D.C. (1997), ISBN 0309056357, <http://www.nap.edu/catalog/10785.html>
- [44] W. Southern: Dealing With Commonwealth Science and Technology Records. Proceedings of the 'Recovering Science' conference, University of Melbourne (1992), <http://www.asap.unimelb.edu.au/confs/recovering/southern.htm>
- [45] CODATA Working Group on Archiving Scientific Data, The Committee on Data for Science and Technology (CODATA) of the International Council for Science, <http://www.nrf.ac.za/codata/>
- [46] Report of the ERPANET / CODATA Workshop "The Selection, Appraisal and Retention of Digital Scientific Data", December 15-17, 2003, Lisbon, Portugal, The Electronic Resource Preservation and Access Network Project and The Committee on Data for Science and Technology (CODATA) of the International Council for Science (2003), <http://www.erpanet.org/www/products/lisbon/LisbonReportFinal.pdf>
- [47] Book of Abstracts of the international Symposium "Ensuring Long-Term Preservation and Adding Value to Scientific and Technical Data", November 5-7, 2002, Centre National d'Etudes Spatiales (CNES), Toulouse, France.
- [48] Bits of Power – Issues in Global Access to Scientific Data. U.S. National Research Council, National Academy Press, Washington, D.C. (1997), ISBN 0309056357, <http://www.nap.edu/catalog/5504.html>
- [49] Report of the Workshop "Interfaces to Scientific Data Archives", March 25-27, 1998, California Institute of Technology, <http://www.cacr.caltech.edu/SDA/ISDA398/isda398.html>
- [50] Squirrel SQL Client is an interesting Open Source project: A Java front end to interact with heterogeneous database environments. It uses the SQL validation web service provided by Mimer Information Technology (cf. Ref. [34]). General description of Squirrel at <http://www-106.ibm.com/developerworks/db2/library/techarticle/dm-0312bhogal/>, download at <http://squirrel-sql.sourceforge.net/>. We have not yet tested this product.
- [51] Two examples are the Access to Archival Databases system (<http://www.archives.gov/aad>) of the U.S. National Archives & Records Administration, and the National Digital Archive of Data Sets (<http://ndad.ulcc.ac.uk>) of the British National Archives.
- [52] The database "FLUGLE" (which we use as an example for demonstration purposes in our figures) is a modified version of the FLUGLE database that comes with the book "Teach Yourself Oracle8 Database Development in 21 Days" by David Lockman, SAMS (1997), ISBN 0672310783 (german edition: Oracle 8 Datenbankentwicklung in 21 Tagen m. CD-ROM, David Lockman, SAMS 1998, ISBN 3-8272-2017-3).
- [53] Mozilla Application Suite and Firefox Web Browser: The Mozilla Foundation, <http://www.mozilla.org>.
- [54] Trivadis (Switzerland) AG, Bern, www.trivadis.com. Trivadis' SIARD project outline can be downloaded: http://www.trivadis.com/Images/bundesarchiv_en_tcm9-849.pdf (english), http://www.trivadis.com/Images/bundesarchiv_de_tcm10-845.pdf (german)
- [55] The Eclipse Project, <http://eclipse.org/>
- [56] The Java Platform, Sun Microsystems. General: <http://www.sun.com/software/learnabout/java/>; Java Foundation Classes (JFC/Swing): <http://java.sun.com/products/jfc/>; Java Database connectivity (JDBC), <http://java.sun.com/products/jdbc/>
- [57] XML Java API for XML Processing (JAXP), <http://java.sun.com/xml/jaxp/>
- [58] XSL Transformations (XSLT) Version 1.0, The Extensible Stylesheet Language Family (XSL), World Wide Web Consortium Recommendation (1999), <http://www.w3.org/TR/xslt>
- [59] ISO/IEC 9075:1999 – 1999 edition of the SQL standard, see Ref. [20]
- [60] ISO/IEC 14977:1996 – Information technology – Syntactic metalanguage – Extended BNF. International Organization for Standardization, and International Electrotechnical Commission. Geneva, Switzerland (1996). The BNF of SQL is described in the SQL standard, see Ref. [59].

A Planetary Data System for the 2006 Mars Reconnaissance Orbiter Era

PV-2004, "Ensuring the Long-Term Preservation and Adding Value to the Scientific and Technical Data"

5-7 October 2004, ESA/ESRIN, Frascati, Italy

J. Steven Hughes⁽¹⁾, Daniel Crichton⁽¹⁾, Gerald Crichton⁽¹⁾,
Ronald Joyner⁽¹⁾, Sean Kelly⁽¹⁾, Chris Mattmann⁽¹⁾, Joel Wilf⁽¹⁾

⁽¹⁾ *Jet Propulsion Laboratory Pasadena
4800 Oak Grove Drive*

Pasadena, CA 91109 USA

*Email: {steve.hughes, dan.crichton, jerry.crichton, ron.joyner,
sean.kelly, chris.mattmann, joel.wilf}@jpl.nasa.gov*

INTRODUCTION

The Planetary Data System (PDS) is the official science data archive for NASA's planetary science community and currently contains about 10 terabytes of data collected from over thirty years of solar system exploration missions. The PDS is an online archive that consists of several geographically distributed science discipline nodes. These nodes support the production of archive quality data products, curate the products, and provide supporting science and technical expertise to the planetary science community.

The success of the PDS is primarily due to the early development and use of a data architecture and data standards for organizing and describing the data in the archive. After being validated against the standards and peer-reviewed by discipline scientists, the data was distributed on physical media to all subscribed users. These "self-contained" archive quality volumes contained not only scientifically useful data products but supporting ancillary data files and informative metadata. The ability of individual scientists to build their own local data library has resulted in high expectations with regard to data availability.

The advent of the Internet combined with sky rocketing costs for distributing physical distribution media forced a review of the distribution process and resulted in the successful development of a middleware-based framework [2] for the on-line distribution of the 2001 Mars Odyssey mission data products. As a result, all missions slated to archive their data are now expecting to use the new distribution capabilities. In addition, it has been recognized that the framework can be used to integrate and simplify the loosely coupled data product production "pipeline" currently used to produce science data products.

The key development challenges now involve software reuse and system deployment. As missions produce larger volumes of data and more product types, the tendency now is to create temporary data nodes at the instrument team institutions to produce and curate the data during the mission. This trend suggests that customized data production and distribution subsystems be configured for specific product types and deployed as needed. The mandate remains however that all data repositories conform to a single data architectural standard and that are capable of being viewed as an integrated whole by external users.

Scalability also poses to be a key challenge for future planetary missions. The science data archived from all planetary exploration missions prior to the 2001 Mars Odyssey mission totals approximately 5 terabytes. The currently active Mars Odyssey mission is expected to double this volume and the 2006 Mars Reconnaissance Orbiter (MRO) is expected to increase the resulting volume by a factor of 10. Future missions such as the 2012 Jupiter Icy Moons Orbiter with its relatively unlimited power supply are predicting even far more data.

This paper will describe several efforts now proposed to address this challenge while providing reuseable, customized, scalable, and remotely deployable software packages that will meet the data intensive requirements of the MRO era and beyond.

DATA SYSTEM REQUIREMENTS

The conversion of the PDS to an online data system necessitated the development of an updated set of system requirements. This was accomplished by a PDS-wide gathering and definition effort that included lessons learned from implementing the 2001 Mars Odyssey distribution system and needs from the project and the science user communities. This input was drafted into a set of requirements by a PDS working group and then edited and documented by the central node systems engineering team.

The resulting Software Requirements Document (SRD) [7] includes a system description, functional requirements, non-functional requirements (system characteristics and constraints), priorities, and an approach for changing the requirements. Subsequent review also partitioned the requirements into those for a "core" system and an "extended" system. The core system includes capabilities for data ingest, validate, track, notify, search, display, retrieve, transfer, and administer. These capabilities are fundamental and were already largely well-defined. The extended system extends the capabilities of the core system with correlative search, advanced display and retrieval, data mining, geometric processing, etc.

DATA DISTRIBUTION –SEARCH, DISPLAY, RETRIEVE, AND NOTIFY

Enabling a user to find, view, and get data and related resources are addressed by the following key functional requirements. A user shall be able to 1) find data based on any indexed attribute, 2) view a visual representations of data through a web interface, 3) download selected products along with associated products, 3) download a subset of selected products, and subscribe to be notified when new data become available.

As previously mentioned, the PDS developed a middleware-based framework for the on-line distribution of the 2001 Mars Odyssey mission data products, called PDSD (see Fig. 1). The deployed data subsystem met the data distribution requirements for the Mars Odyssey mission, is now successfully distributing almost all the legacy data in the archive, and is expected to meet the distribution requirements for missions planned through the MRO era.




Figure 1 - PDS-D Configuration for Mars Odyssey

PDS-D is based on a four-tiered information architecture that consists of application, web, service, and storage tiers. The multi-tiered architecture partitions the system into separately managed components and also provides a means by which the system can be deployed with minimal impact on existing resources. For example, the Planetary Atlas developed by the PDS imaging node is the primary application for searching and retrieving imaging products from the archive and represented a significant investment of developmental resources. The Planetary Atlas was integrated into the client tier of the architecture as can be seen in Fig. 1. The service (middleware) and storage (repositories) tiers are also illustrated in the figure.

In the storage tier, the data archive remains geographically distributed and locally managed by the science domain experts. The development of this tier leveraged heavily on the data architecture previously created by the PDS for the archive. Organized into archive volumes on optical media, the data with its accompanying metadata and ancillary data was copied from the archival volumes onto mass storage devices at a local node to make the collection accessible through the service tier.

The service tier includes the middleware as well as service components that support the search, retrieval, and transformation of data from the geographically distributed components of the storage tier. Key service components include product servers that provide a common interface for the retrieval and transformation of data products from a data repository. Profile servers provide a common interface for product search. A query service manages queries entering the system by broadcasting them to the appropriate product or profile servers and compiling the results for transmission back to the user. The service tier and its implementation using the Object-Oriented Data Technology (OODT) platform is described in more detail in [1,2,3].

Application developers at the client tier can interface directly with the service tier through JAVA APIs. However the web tier includes a web server to allow http level access. In fact the majority of applications developed for the infrastructure have used http level access for data search and retrieval.

A key design principle that has evolved out of the PDS experience is the importance of separating the data architecture from the technology architecture [4]. This principal naturally arises by observing the differences in the life cycles of data architectures versus technology architectures. Data architectures typically model a specific domain and evolve as the domain evolves, typically very slowly. However the technologies used to implement systems evolve much more rapidly and are often out-of-date soon after an implementation has finally been deployed. The independence of the two architectures allows both to evolve as needed to meet the needs of the enterprise. The OODT framework uses XML to manage the metadata both for resource descriptions and messaging and the ISO/IEC 11179 [6] specification as a standard for vocabularies. A novel XML document called a resource profile is used to provide a single uniform view of all resource descriptions for both querying and returned results [1].

The data architecture and data standards together act as important leverage for the development and implementation of the technology architecture. For example, the existence of and adherence to a data architecture not only made the configuration of the PDS storage tier very easy, but the wealth of metadata in the archive also makes the indexing of products for catalog search relatively easy. In addition the metadata also provides a rich source of information for data mining and correlative search applications.

DATA INGESTION – INGEST, VALIDATE, AND TRACK

Governing the incorporation of mission archive data into the system is addressed by the following key functional requirements. The system shall a) receive data, b) catalog all metadata or “data descriptions”, c) validate the data delivery against a specification to ensure that it is compliant with the data standard and d) keep track of the status of every data delivery. These core functional requirements are not much effected by specific requirements from the upcoming MRO and Cassini missions. However MRO-specific requirements do raise data volume issues and Cassini-specific requirements raise model complexity issues associated with the large number of data types expected from the mission. Two key issues being addressed are 1) the validation functional component needs to scale to meet the combined load of the two missions and 2) even though the data will be available online, the requested distribution volume to the planetary community will be very difficult to meet using commonly available web technologies.

The PDS is currently in the process of developing the data ingestion, validation, and tracking subsystem to meet the given requirements. The multitiered architecture based on the OODT middleware that was introduced for the PDS-D implementation will again provide the base platform for development. This platform has matured and is now “open source” and a “web service” interface has been enabled. The PDS data repository structure introduced for PDS-D will provide the data storage component.

Key to the validation requirement is the existence of a specification against which all archived data are validated. This is embodied in the PDS data standards as documented in the PDS standards reference. This reference document is being edited to make it more suitable as a specification. After validation the results are used as a criteria for accepting a delivery for ingestion. Levels of validation include syntactic or checking that the documents are “well formed”, semantic, or checking that the data is appropriately described using domain knowledge, and quality checks that ensure the delivery is complete, consistent, and useable.

The production of archive quality data products requires a well-defined process and “work flow” that manages the movement of individual data product through a production “pipeline”. Key functions include controlling product versioning, state changes, and the application of either automated or manual process. An Ingest Workflow Manager is planned, based on the successful ground data system built for the SeaWinds project. This system, built from an existing component of the OODT framework called the Catalog and Archive Service (CAS), will now be validated in a multi-mission environment as shown in Fig. 2. The CAS provides a rule-based infrastructure focused on tying together a set of rules used to construct both a catalog and a repository into the workflow. The rules are implemented as set of Java-based processes that run as part of the process of ingesting a data product. The ingestion process is part of a larger transaction which means that products that fail to be properly ingested will cause the CAS to rollback. The multi-mission configuration is supported by enabling highly diverse processing scenarios. These include scenarios from simple validation to complex processing and creation of new data products. The CAS will extend the distributed PDS by enabling access to the catalog and repository using the OODT information architecture which provides common definitions for data queries and product resources.




Figure 2 - PDS Distributed Ingest Architecture

Similarly a prototype Data Tracker developed for the Cassini mission will be generalized to meet multi-mission data tracking requirements. It should be noted that the validation requirements just discussed focus on the data production pipeline. PDS policy still requires passing a science peer review before a data collection can be considered archived. Fig. 3 illustrates a mission data flow into the PDS for distribution and archive. It abstracts the multi-tiered distribution architecture of Fig. 1 into a single component for a single node and integrates the ingestion, validation, and tracking elements into an additional component.




Figure 3 - Product Data Flow with Ingestion, Validation, and Distribution

CONCLUSION

The Planetary Data System archives data for the planetary science community. Although the total data volume in the archive is not large relative to other science domains such as Earth Science, the planetary science domain is much more complex than most other space science domains since they involve dynamic systems that include orbiting target bodies, moving instrument platforms, complex instruments, a plethora of frame of references, changing sources of light, etc. The development of a data architecture and data standards in the early years of the PDS enabled the creation of a data archive consistent in its structure, meaning, and organization as well as rich in descriptive information. The advent of the web has provided the technologies needed to make the planetary science archive available to a wider range of customers in increasingly more useful and sophisticated ways.

Movement towards standard mechanisms to catalog, validate, process and distribute data via distributed software interfaces is enabling the PDS to handle the increase in volume and complexity for future missions. Key to this advancement is the formulation of an information architecture that provides for independent data and technology architectures and the development of a multi-tiered infrastructure that allows the integration of heterogeneous distributed data repositories into single system with common system interfaces. [1,2,3] The planned implementation of the ingest/validate/track functional components into the existing infrastructure will provide a highly-automated end-to-end data production and distribution system for the planetary science community. This is a long awaited goal that few other science domains of this breadth have accomplished.

Future development includes addressing the desires of some MRO scientists to distribute entire collections of highly derived data to thousands of users at the end of the mission. Calculations based on current estimates predict this would necessitate the distribution of petabytes of data in the period of a few months. A proposal has been written to research available technologies such as grid services [5] and to deploy a prototype infrastructure across the planetary science community.

REFERENCES

- [1] Crichton D, Hughes JS, Hyon J, Kelly S., "Science Search and Retrieval using XML". Proceedings of the 2nd National Conference on Scientific and Technical Data, National Academy of Science, Washington D.C, 2000.
- [2] Kelly S, Crichton D, Hughes JS., "Deploying Object Oriented Data Technology to the Planetary Data System", Proceedings of the 34th Lunar and Planetary Science Conference 1607, 2003.
- [3] Mattmann C, Ramirez P, Crichton D, Hughes, JS, "Packaging Data Products using Data Grid Middleware for Deep Space Mission Systems", Proceedings of the 8th International Conference on Space Operations (Space-Ops2004), 2004, in press.
- [4] Consultative Committee on Space Data Systems, "Space Information Architecture", White Paper, Information Architecture Working Group. February 2004, in press.
- [5] Foster I, Kesselman C, Nick J, Tuecke S., "The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration", Open Grid Service Infrastructure WG, Global Grid Forum, 2002.
- [6] ISO/IEC 1999. Framework for the Specification and Standardization of Data Elements 11179-1, Specification and Standardization of Data Elements 11179. International Organization For Standardization.
- [7] The Planetary Data System- System Requirements Document (SRD), Version 1.2, JPL Internal Document, November 2003.

The UK Digital Curation Centre

5-7 October 2004, ESA/ESRIN, Frascati, Italy

David Giarretta⁽¹⁾, Peter Burnhill⁽²⁾, Seamus Ross⁽³⁾
Liz Lyon⁽⁴⁾, Peter Buneman⁽²⁾

⁽¹⁾Digital Curation Centre and
CCLRC, Rutherford Appleton Laboratory, Chilton, Didcot, Oxon OX11 0QX, UK
Email: d.l.giarretta@rl.ac.uk

⁽²⁾Digital Curation Centre and
Edinburgh University, Edinburgh, EH8 9LJ, Scotland
Email: P.Burnhill@ed.ac.uk

⁽³⁾Digital Curation Centre and
HATII, George Service House, University of Glasgow, G12 8QQ, UK
Email: s.ross@hatii.arts.gla.ac.uk

⁽⁴⁾Digital Curation Centre and
UKOLN, University of Bath, Bath, BA2 7AY, UK
Email: e.j.lyon@ukoln.ac.uk

INTRODUCTION

The UK has set up the Digital Curation Centre (DCC). The aims of the DCC (<http://www.dcc.ac.uk>) include

- 1) Establishing a vibrant research programme - addressing wider issues of data curation
- 2) Nurturing strong community relationships - forming and extending an Associates Network, engaging with scientific digital curators
- 3) Developing services - for testing and evaluating tools, methods, standards and policies in realistic settings - offering a repository of tools and technical information, a focal point for digital curators
- 4) Achieving the 'virtuous circle' - feeding expertise, experience and need into its research programme on data curation - transforming research-led innovation into services that enhance productivity of practice.

This paper will describe some aspects of the DCC in some detail. It will give some idea of the basis on which this unique organisation will supply advice, tools and underpinning research to support data curation for scientific as well as other scholarly data.

Definition of Digital Curation

The term Digital Curation is a rather recent invention. The e-Science Curation Report (2003) by Lord and Macdonald (http://www.jisc.ac.uk/uploaded_documents/e-ScienceReportFinal.pdf) proposed the following distinctions:

- **Curation** : The activity of, managing and promoting the use of data from its point of creation, to ensure it is fit for contemporary purpose, and available for discovery and re-use. For dynamic datasets this may mean continuous enrichment or updating to keep it fit for purpose. Higher levels of curation will also involve maintaining links with annotation and with other published materials.
- **Archiving** : A curation activity which ensures that data is properly selected, stored, can be accessed and that its logical and physical integrity is maintained over time, including security and authenticity.
- **Preservation** : An activity within archiving in which specific items of data are maintained over time so that they can still be accessed and understood through changes in technology.

There does seem to be a lack of clarity, and this is at least partly because there are a number of terms used singly or in combination which have more or less clear definition, but they are used without adequate care. The danger is that the reader (or listener) may believe that something extra is implied. Some terms that are worth distinguishing are:

- **data preservation** : a general term probably equivalent to digital preservation in this context
- **digital preservation** : could be, and probably is, often interpreted as simply ensuring the original bits and bytes are accessible

- **digital information preservation** : this is what is referred to in the OAIS standard - what is important is not the original "bits and bytes" but the content. An OAIS ensures that the content is accessible, understandable and usable.
- **curation** : general term - taking care of things
- **data curation** : looking after and adding value to data
- **digital curation** : looking after and somehow "adding value" to digital data. This probably implies creating some new data from the existing, in order to make the latter more useful and "fit for purpose".
- **information curation** : not seen in the wild

FOUNDATIONS

The DCC uses the OAIS Reference Model (ISO 14721, <http://ccsds.org/documents/650x0b1.pdf>) and its view of information preservation as the basis for its preservation work, but then adds in ideas from many other sources in order to cover areas which are not within the remit of OAIS. This starting point is a pragmatic one: OAIS is recognised as a significant standard in this area.

In addition to information preservation, the term Curation has been used to indicate support for current research activities using that information, including using information in new ways and also publishing results based on the information. It is perfectly possible for, and there are many examples of, digital repositories supporting current research not to have any long term preservation aspirations. Similarly it is possible to have a focus on preservation without much regard for supporting current research. However neither of these would be the right course for the DCC.

There are very many existing projects to support current research activities including many e-Science projects and the JISC Information Environment Architecture (<http://www.ukoln.ac.uk/distributed-systems/jisc-ie/arch/>) - all of which have particular relevance to the DCC. It does not seem sensible to duplicate or re-invent these. However we argue below that an architecture which tackles preservation issues – ensuring that information is usable in the future where users are unfamiliar with the data – will also be significant in supporting current usage, especially where, again, users are unfamiliar with the data.

We therefore take the view that we should be guided by long-term preservation aspects, and try to ensure that components of the preservation architecture can supplement other “current use” architectures. To promote this we emphasise “interoperability” and “automated use” as far as possible.

By “interoperability” we mean here the ability of separate systems (possibly covering different disciplines and built using different architectures) to exchange and use each others’ data. This is important because we cannot dictate information architectures for “current use” being developed now, much less those developed in the future. The term “automated use” denotes the desire for systems to be able to deal with information without the need for human intervention, and in particular without the need for a human to read and interpret documentation associated with the information – especially important as we deal with increasing amounts of information from more and more sources.

STRUCTURE OF THE DIGITAL CURATION CENTRE

The Digital Curation Centre is a consortium consisting of the Universities of Edinburgh, Glasgow, Bath and CCLRC. As shown in Fig. 1 we hope to take advantage of a “virtuous circle” between Community Support and Outreach, Research, Development and Services. In addition the DCC will work with industry, standards bodies and many other collaborators. The following sections give a flavour of our thinking at the time of writing as we initiate our first phase.




Figure 1 Structure of the Digital Curation Centre

PRESERVATION INFRASTRUCTURE

The OAIS Reference Model has many aspects, including an Information Model, a Functional Model, and an Information Flow model. All these play their part. However two key concepts drive this document (1) the Designated Community and its Knowledge Base and (2) Representation Net, defined as follows.

Designated Community is defined as an identified group of potential Consumers who should be able to understand a particular set of information. The Designated Community may be composed of multiple user communities. **Knowledge Base** is defined as a set of information, incorporated by a person or system, that allows that person or system to understand received information.

OAIS: Representation Network

A basic concept of the OAIS Reference Model (ISO 14721, <http://www.ccsds.org/documents/650x0b1.pdf>) is that of information being a combination of data and Representation Information. The UML diagram in Fig. 2 illustrates this concept. The Information Object is composed of a Data Object that is either physical or digital, and the Representation Information that allows for the full interpretation of the data into meaningful information. This model is valid for all the types of information in an OAIS.




Figure 2 OAIS Information Object




Figure 3 OAIS Representation Information Object

The UML diagram in Fig. 2 means that

- an **Information Object** is made up of a **Data Object** and **Representation Information**
- A **Data Object** can be either a **Physical Object** or a **Digital Object** . An example of the former is a piece of paper or a rock sample.
- A **Digital Object** is made up of one or more **Bits** .
- A **Data Object** is interpreted using **Representation Information**
- **Representation Information** is itself interpreted using further **Representation Information**

Fig. 2 shows that Representation Information may contain references to other Representation Information. When this is coupled with the fact that Representation Information is an Information Object that may have its own Digital Object and other Representation Information associated with understanding each Digital Object, as shown in a compact form by the *interpreted using*. association, the resulting set of objects can be referred to as a Representation Network. The Representation component Fig. 3 shows more details and in particular breaks out the semantic and structural information as well as recognising that there may be “Other” representation information such as software.

The recursion of the Representation Information will ultimately stop at a physical object such as a printed document (ISO standard, informal standard, notes, publications etc) but use of things like paper documentation would tend to prevent "automated use" and "interoperability", and also complete resolution of the complete Representation Network to this level would be an almost impossible task. Therefore we would prefer to stop earlier. In particular we can stop for a particular Designated Community when the Representation Information can be understood with that Designated Community’s Knowledge Base.

For example a science file in FITS format would be directly understood by someone who knew how to handle this format – someone whose Knowledge Base includes FITS – for example has some appropriate software. Someone whose Knowledge Base does not include FITS would need additional Representation Information – for example would have to be provided with some software or the written FITS standard.

FUNDAMENTAL ISSUES OF PRESERVATION AND USE

Given a file or a stream of bits how does one know what Representation Information is needed? (this question applies to Representation Information itself as well as to the digital objects we are primarily interested in preserving and using); how does one know, for example, if this thing is in FITS format?

1. Someone may simply “know” what it is and how to deal with it i.e. the bits are within the Knowledge Base
2. One may have an associated label which points to the appropriate Representation Information.
3. One may be able to recognise the format by looking for various types of patterns.
4. One may feed the bits into all available interpreters to see which accept the data as valid
5. Other means....

Of the above, if (1) does not apply then only (2) is reliable because the others rely on some form or other of pattern recognition and there is no guarantee that any pattern is unique. Even if the File Format is unique the possible associated semantics will almost certainly not be so.

However if no label is available then one of the other methods must be used, as would be the case for data rescue (in the sense of data inherited without adequate metadata, but not itself corrupt).

The implications for a preservation infrastructure are that a label must be attached to each piece of digital object as a necessary (but not sufficient) condition for long-term preservation – note that this is some kind of logical attachment or packaging defined, for example, by the DCC or by other means. The label should at least identify Representation Information. For long-term preservation this label must therefore be a DCC persistent identifier. Note that in order to allow some normalisation e.g. in the case of a compressed tar file containing a FITS file we may wish to identify “compress”, “tar” and “FITS” separately, or identify variety of combinations of syntax and semantics. To allow for this the label may have some structure itself – in fact it may itself be a digital object.....however we would probably want to prevent too much diversity. On the other hand we would probably want to cope with a variety of labelling since we would need to support a variety of standards.

All the above considerations apply equally to the use of digital information separated from the original producers either by time or by “knowledge base”, and hence a solution should be useful both for long-term preservation as well as “current use” of digital information.

In order for the Representation Information to be persistent then it should either be held with the data object itself or be part of a central repository – part of the DCC. Thus the DCC needs a DCC Representation Information Repository. Because the long-term curation of this Representation Information would have to be guaranteed, adequate succession planning would have to be put in place, for example with a body of guaranteed longevity. In fact we would hope that the Representation Information Repository would develop into a distributed, global, collaboration.

This (global, distributed, persistent) repository would include

- a Format Repository (covering structural information) - automated use would be supported by use of formal description languages such as EAST (ISO 15889, <http://east.cnes.fr/>) or DFDL (<http://forge.gridforum.org/projects/dfdl-wg/>). A backstop would be human readable documents such as the appropriate ISO, or other, standard underpinning the format - for example the FITS standard documents, the ISO 9660 standard etc
- a Semantic Repository with, for example, Data Dictionaries and Ontologies - a backstop would be human readable documents such as code books, human language grammar books and dictionaries etc
- Software Repository – with appropriate emulation capabilities

Each piece of digital Representation Information is also a digital object – which is understood either by the users’ Knowledge Base OR by further Representation Information. Therefore each piece of Representation Information also has a label pointing to further Representation Information.

At any particular time the Representation Network for a given digital object need not be complete – it can be terminated at a point determined by the Knowledge Base – but *which* Knowledge Base? The answer to this comes in the next section.




Figure 4 Representation Information (RI) Architecture

OAIS: Designated Community

A Designated Community has, at any particular time, a particular **Knowledge Base**. For a specific Designated Community this Knowledge Base will evolve with time; in addition the definition of the appropriate Designated Community for a dataset may be changed. The importance of identifying the Designated Community for a data object or more likely a collection in an archive is that it allows the archive to limit the amount of Representation Information required for any particular digital object. Without this limitation the archive would, in principle, have the impossible task of collecting all possible Representation Information. Techniques must be created for defining a Knowledge Base, linking a Knowledge Base to a Designated Community and linking Representation Information to a Knowledge Base if possible.

TYPES OF REPRESENTATION INFORMATION

Structure – including Formats

We believe that it is useful to distinguish between

- formats which are used mainly for rendering – to be used by a human being rather than some computer process, and
- formats used for automated processing

The former include many commercially based formats such as the succession of wordprocessor formats. The details of current commercial formats are likely to be proprietary and difficult or impossible to obtain. On the other hand the format is more likely to be available when that format is no longer the “current” version – which is when we would actually need it.

The latter are more likely to be simpler, with Open Source access software and more easily independently describable. It is proposed that we focus initially on this latter set of formats – although we should not neglect the former type of formats, at least to the level of collecting information about them - enough to start the associated Representation Net for each format.

There are tools, noted above, which can describe digital information in a way suitable for automated processing. EAST and DFDL currently give access to individual components such as numbers or arrays of numbers within a particular format. It may be useful to define useful scientific objects in order to facilitate automated processing. Data virtualisation techniques being developed in the Data Grid are likely to be useful.

In a similar way it may be useful to define some “humanities objects” for the same reason. Some of these, for example images – simple or multispectral – might be the same as their scientific counterparts. Others might be completely new, for example virtual reality city scapes to capture archaeological surveys.

PRONOM (<http://www.records.pro.gov.uk/pronom/>) and the Global Digital Format Registry (GDFR <http://hul.harvard.edu/gdfr/>) both aim to provide information about file formats - although they both focus on “office” type formats. They each have their own data model and plans for service delivery. Neither seem to put their plans into the wider context of complete Representation Information. A trial GDFR implementation provides some transformation services based on the Typed Object Model (TOM <http://tom.library.upenn.edu/>). Perhaps the most functional project is Presidential Electronic Records Project Operational System (PERPOS); they have collected 250+ legacy formats recognised by magic numbers and file structure rather than file extensions and has associated viewers and aim to provide services for file format recognition and transformation. The DCC will try to work with all these projects to build a global representation information repository.

Semantics

In principle we could plunge into developing tools for Ontologies. This is an active area of research by many groups and we are not likely to make rapid progress, although a survey of the current work would be in order. A more tractable topic (important from a pragmatic point of view) to tackle would be the simpler one of Data Dictionaries. A number of ISO standards exist already (ISO 11179 and the set consisting of ISO 21961, 21962 & 22643 - <http://www.ccsds.org/documents/647x3b1.pdf>).

Time Dependent Information

Many, some would say most, datasets change over time and the state at each particular moment in time may be important. This is an important area requiring further research, however from the point of view in this document it may be useful to break the issue into separate parts.

- at each moment in time we could, in principle, take a snapshot and store it. That snapshot has its associated Representation Net.
- efficient storage of a series of snapshots may lead one to store differences or include time tags in the data (see for example *Peter Buneman, Sanjeev Khanna, and Wang-Chiew Tan. On the Propagation of Deletions and Annotations through Views. In Proceedings of 21st ACM Symposium on Principles of Database Systems.*). Additional Representation Information would be needed which describes how to get to a particular time's snapshot from the efficiently encoded version.

Annotation and other time dependent metadata

Annotation was identified in our bid as an area in itself requiring significant research effort. Insofar as annotations and other metadata are themselves data, the considerations of time dependence noted above apply also. Within the OAIS Information Model, Annotation, Provenance etc are part of Preservation Description Information which is discussed in more detail below.

These are area of active research within the consortium and the DCC should be able to provide

- advice and well tested tools for certain forms of efficient encoding of time dependent information
- advice on annotation
- identifiers and Representation, perhaps in the form of software, for the associated encodings

Actions and Processes

Some information has, as an integral part of its content, an implicit or explicit process associated with it – this could be argued to be a type of semantics, however it is probably sufficiently different to need special classification. Examples of this include databases or other time dependent or reactive systems such as Neural Nets.

The process may be implicitly encoded in the data, for example with the scheme for encoding time dependence in XML data as noted above. Alternatively the process may be held in the Representation Information – possibly as software. Amongst many other possibilities under this topic, Software and Software Emulation are among the most interesting (<http://www.dlib.org/dlib/october00/granger/10granger.html>).

It may be possible to develop a Universal Virtual Computer (UVC) as outlined by Lorie (<http://www.rlg.org/preserv/diginews/diginews5-3.html#feature2>). However, recognising that one of the prime desirable features of a UVC is that it is well defined and can be implemented on numerous architectures, it may be possible to use something already in place, namely the JAVA Virtual Machine (JVM, <http://java.sun.com/docs/books/vmspec/>).

Persistent IDs

A number of persistent IDs are available including DoI (<http://www.doi.org/>), CCSDS Unique ID (<http://www.ccsds.org/documents/A31x0y1.pdf>) which uses *ISO/IEC 6523-1:1998 Information technology. Structure for the identification of organizations and organization parts. Part 1: Identification of organization identification schemes, 1998.* and *ISO/IEC 8824-1:1998 Information technology. Abstract Syntax Notation One (ASN.1): Specification of basic notation, 1998.* The DCC could sign up to something like the DoI or could set up an independent system. As for the long-term repository to which the Persistent IDs point, that could initially be provided by the DCC itself, however succession planning is needed and agreement should be reached with an organisation of guaranteed long-term existence.

ARCHIVAL INFORMATION PACKAGE

An Archival Information Package (AIP) is defined to provide a concise way of referring to a set of information that has, in principle, all the qualities needed for permanent, or indefinite, Long Term Preservation of a designated Information Object. Since the AIP is ready for long-term preservation it can be used as the basis for defining several important structures.

The AIP is a logical definition. For practical use an encoding must be provided. METS packaging (<http://www.loc.gov/standards/mets/>) is an XML Schema which is based on the AIP. However concern has been expressed about weakness in the way the Representation Information is dealt with by METS, and other schemes are being investigated (<http://www.ccsds.org>)




Figure 5 OAIS full Archival Information Package

METADATA FOR PRESERVATION

OAIS defines Preservation Description Information as:




Figure 6 OAIS Preservation Description Information

Starting with the AIP components OCLC has put together a definition of metadata for preservation (http://www.oclc.org/research/projects/pmwg/pm_framework.pdf). This work could form the starting point for a DCC definition of preservation metadata, taking into account the additional components in these diagrams. Of particular importance are the following.

Reference Information

Reference Information is *the information that identifies, and if necessary describes, one or more mechanisms used to provide assigned identifiers for the Content Information. It also provides identifiers that allow outside systems to refer, unambiguously, to a particular Content Information. An example of Reference Information is an ISBN.*

The Persistent ID discussed above could fulfill this role.

Provenance Information

Provenance Information is *the information that documents the history of the Content Information. This information tells the origin or source of the Content Information, any changes that may have taken place since it was originated, and who has had custody of it since it was originated. Examples of Provenance Information are the principal investigator who recorded the data, and the information concerning its storage, handling, and migration.*

This is an active area of research, within the DCC.

Context Information

Context Information is *the information that documents the relationships of the Content Information to its environment. This includes why the Content Information was created and how it relates to other Content Information objects.*

Fixity Information

Fixity Information is defined as *the information which documents the authentication mechanisms and provides authentication keys to ensure that the Content Information object has not been altered in an undocumented manner. An example is a Cyclical Redundancy Check (CRC) code for a file.* Experience has shown that even with a single computer system, error-free transfers of data cannot be taken for granted. The Atlas Petabyte Store at CCLRC uses a sophisticated checksum regime in order to confirm the correct transcription of data. No doubt other data stores have similar systems in place. The DCC could compare the techniques used and recommend a standard checking regime.

It may be that additional checks may be required and a full authentication system is required. The DCC could, for example, recommend a public/private key system and keep the keys stored in perpetuity in order for authenticity to be confirmed. Over the long term increasingly sophisticated systems are likely to be required and an associated process defined for strengthening authentication methods.

CERTIFICATION

In order to have a certification process there needs to be in place a Standard (preferably one or more ISO standards) against which to do the certification. The RLG Digital Certification Task Force (http://www.rlg.org/en/page.php?Page_ID=367) is working to produce a certification process. The documents produced will probably become ISO standards via CCSDS (the body which produced the OAIS Reference Model). Indeed this work can be seen as a part of a suite of OAIS-related standards.

Following the production of the standard, accreditation bodies and certification bodies would need to be set up. Contact would have to be made with an overall body such as the International Accreditation Forum (IAF <http://www.iaf.nu/>) or International Laboratory Accreditation Cooperation (ILAC <http://www.ilac.org/>) or a regional body such as the European cooperation for Accreditation (<http://www.european-accreditation.org/>). Members of the DCC are involved in these efforts and the DCC should play a significant role in this.

CONCLUSIONS

The remit of the UK's Digital Curation Centre to cover both science as well as the humanities and digital libraries gives it a special view on digital curation which we believe will allow the DCC to be a major force in the arena within the UK and internationally. This paper has outlined the basis of some of our plans in what is, at the time of writing, our start-up phase, and the DCC will develop a number of "virtual circles" over next few years, building on this initial phase. There is in addition activity in the areas of Outreach and User Requirements, Services and Research which are not described here. Up to date information should be available on the DCC web site <http://www.dcc.ac.uk>.

Digital Libraries and Grid Technologies as Infrastructure for Earth Observation Data Archives Exploitation and Long-Term Preservation

**Authors: Luigi Fusco⁽¹⁾, Joost van Bemmelen⁽²⁾,
Veronica Guidetti⁽¹⁾, Donatella Castelli⁽³⁾**

*⁽¹⁾ ESA ESRIN (Italy)
Via Galileo Galilei. 00044 Frascati (RM)
Email: luigi.fusco@esa.int, veronica.guidetti@esa.int*

*⁽²⁾ Intecs c/o ESA-ESRIN (Italy)
Via Galileo Galilei. 00044 Frascati (RM)
Email: Joost.van.Bemmelen@esa.int*

*⁽³⁾ Consiglio Nazionale delle Ricerche - Institute of information science and tech. (Italy)
Via Giuseppe Moruzzi, 1 - 56124 Pisa PI Toscana
Email: donatella.castelli@isti.cnr.it*

Abstract

Among other tasks, the European Space Agency ESRIN establishment is in charge of the operations of the reference European Earth Observation (EO) payload data exploitation infrastructure and, in particular, of the facilities dealing with the Envisat and Ers missions (ERS-1 launched in '91 and retired from operation in early 2000, and ERS-2 launched in '96). The historical ESA EO archives account for Petabytes data holding, which is augmented, since the launch of Envisat in 2002, by some 500 Terabytes per year.

The EO community accounts for thousands of users dislocated across Europe and the rest of the world and for distributed management of the historical and present data holdings. Emerging institutional and international environmental initiatives (e.g., GMES, Global Monitoring for Environment and Security) need complete appraisal of full historical collections of uniform and objective measurements, including the performed data elaboration, scientific analysis, models and results.

The access to and utilisation of historical archives is an important measurement of long-term data preservation. Improvement of user access to real-time and historical EO data holdings is a continuous challenge at programmatic, technological and operational level.

Long-term data preservation was recently (December 2003) discussed at the ERPANET/CODATA workshop on "The Selection, Appraisal and Retention of Digital Scientific Data" with the clear vision that the issue is common to all sciences.

Emerging ICT technologies, such as Digital Libraries (DL) and Grid, which are based on networked infrastructures, are essential components in facilitating the handling of multiple and partial copies of distributed data sets. Therefore DL and Grid can play an important role in long-term data preservation.

The paper highlights the progress made in integrating such technologies to demonstrate the provision of a new generation of data and information access services for the Earth and Environmental Science communities.

Introduction

Earth observation data are large in volume and range from the local to global in context. Within initiatives like the ESA/EC Global Monitoring for Environment and Security (GMES) [1] and ESA' Oxygen [2], large amounts of digital data, that are potentially of great value to research scientists in the Earth observation community and beyond, need to be acquired, processed, distributed, used and archived in various facilities around the world.

For example, Envisat [3], the advanced European polar-orbiting Earth observation satellite, carries a payload of 10 instruments that provide huge amounts of measurements of the atmosphere, ocean, land, and ice. These data are processed, corrected and/or elaborated using auxiliary data, models and other relevant information in support to Earth science researches (e.g., monitoring the evolution of environmental and climatic changes). All of these need to be preserved for future generations to better understand the evolution of our Earth. That is, preservation doesn't relate to the data only, but implies also the maintenance of information regarding when and how such data were processed, the reason(s) why, or why this way. It may include knowledge, documented results, scientific publications, and any other information needed to support scientists in their research.

Currently, however, there is no clear mandate to preserve Earth observation mission data, relevant information and knowledge at the European level and the responsibility falls under the remit of the individual mission owners and/or national archive holders. Coordinating efforts on standards and approaches to preserve the most valuable European Earth observation data will be required to guarantee access and reusability of these, often distributed, data.

It should be noted that the availability and accessibility to most of the abovementioned information is considered relevant for efficient and adequate exploitation of Earth observation data and derived products as well. As such, their preservation is also considered important for optimal short-term archive and data reuse by different actors, including multi-disciplinary exchange of experiences on the same dataset. In other words, problems

encountered in long-term preservation certainly include those encountered in short-term exploitation, and as such the paper's focus is on some aspects of data preservation.

Solid infrastructures are needed to enable timely access, now and in the future. Much more digital content is available and worth preserving; researchers increasingly depend on digital resources and assume that they will be preserved [4].

The following paragraphs describe challenges in long-term data preservation, discuss the digital library and Grid technologies in relation to long-term data preservation, briefly discuss semantic technologies, and provide some abstracts of related projects. At the end, a way forward is proposed.

Challenges in Long-Term Data Preservation

The workshop on "Research Challenges in Digital Archiving and Long-term Preservation" [4] that was held in 2002, identified many challenges in long-term data preservation. The overall challenge is guaranteeing access to and usability of data, independent from underlying hardware and software systems; systems evolve along time, technologies are being renewed and replaced, data formats may change, and so may as well related information and scientific knowledge.

The last few years, different communities have tackled this problem experimenting different technologies. These include Semantic Web focussing on semantic-web access, data grid technology focussing on management of distributed data, digital library technology focussing on publication and persistent archive technology focussing on management of technology evolution [5].

In Earth Science, accessing historical data, information and related knowledge may be nowadays quite complex and sometimes impossible, both, due to the lack of descriptive information (metadata) that could provide the context in which they fit but also because of the lack of the information and knowledge themselves.

Different initiatives are focussing on these issues [6]. One of them is ERPANET / CODATA [7] but also the Persistent Archive Research Group and the Data Format Description Language research group, both part of the Global Grid Forum, are looking at similar questions [5,8,9]. Within the EO-community CEOS, the Committee of Earth Observation Satellites, is looking at the use of XML for Science Data Access [10]. Based on results achieved in ongoing projects and expected results in planned projects, a possible technical solution to approach long-term data preservation may consider technologies as digital libraries and Grid.

Digital Libraries and Grid

Digital libraries and Grid are two quite new technologies that can be used for long-term preservation. In 1993 the term "digital library" (DL) appeared for the first time, designating collections of electronic information that are maintained at and possessed by the library itself. But the concept of "digital library" does not only point to the collection itself: it has to be thought of as a whole bunch of functions and services, including storage, discovery, retrieval, and conservation of the data. From a recent analysis of 65 definitions of DL the following common characteristics were derived [11]:

Digital libraries ideally:

- serve a defined community or set of communities
- are underpinned by a unified and logical organizational structure
- incorporate learning as well as access
- provide fast and efficient access, with multiple access modes
- provide free access (perhaps just to the specified community)
- own and control their resources (some of which may be purchased)
- have collections which
 - are large, and persist over time
 - are well organized and managed
 - contain many formats
 - contain objects, not just representations
 - contain objects which may be otherwise unobtainable
 - contain objects which are digital *ab origine*

DLs are seen as an essential element for communication and collaboration among scientists and represent the meeting point of a large number of disciplines and fields including data management, information retrieval, library sciences, document management, information systems, the web, image processing, artificial intelligence, human-computer interaction etc.

Most of current systems however, run in a single organisation, mainly handling textual documents since the handling and preservation of large sets of multimedia documents and data require computational and storage resources that are rarely available in a single organisation. Furthermore, the automatic elaboration of multimedia data, e.g., the automatic extraction of the contents description of multimedia documents, is often too expensive. Recently, a few experimental DL Management Systems, based on distributed architectures, have been proposed.

The term Grid (sometimes also called the new world wide web) firstly appeared in the nineties, and was used to denote a new distributed computing infrastructure for highly efficient solutions to scientific and engineering applications. At the time of writing, according to Foster [12], Grid is a system that:

- coordinates resources that are not subject to centralized control;
- uses standard, open, general-purpose protocols and interfaces;
- delivers nontrivial qualities of service.

Grid infrastructures have been experimented and implemented throughout the world the last few years. An example is the European Data Grid (EDG) project [13] that was the first large-scale international grid project and the first aiming to deliver a grid infrastructure to several different "Virtual Organisations" (High Energy Physics, experiments, Biology and Earth Observation) simultaneously. Some Grids have been demonstrated in operational environments and on-going work is now aimed at creating a reliable and dependable European Grid Infrastructure (e.g., in the EC project EGEE: Enabling Grids for E-science in Europe [13]).

Grid could be a valuable complementarity technology to the DL as it addresses the major DL architecture requirements, i.e. openness, scalability, security and quality. Moreover, these technologies seem to marry very well to build so-called Virtual Digital Libraries (VDL), i.e. transient DLs based on shared computational, multimedia and multi-type content and application resources. Here, techniques for data replication and security handling developed in the Grid area will contribute greatly to the definition of new DL preservation techniques, also because grids provide the abstraction mechanism needed to deal with heterogeneous hardware and software environments [14]. Furthermore, functions like content feature extraction, summarization, automatic content source description, etc. on video images and sound, which are based on complex and time-consuming algorithms, will become viable with acceptable performance. Intertwining Grid and DL technologies will also lead to developing next generation information networks on the basis of Grid technology that can be applied in adapted in many different domains.

In summary, the long-term data preservation has to be based on a distributed environment capable to handle multiple copies of the same information. Grid and DL technology could help in performing long-term data preservation since, as said above, the preservation task of migrating from old to new technology is really similar to manage access to data distributed across multiple sites while the organization of data and metadata in information collections requires discovery and access techniques as provided within DLs.

Semantic Technologies

Another group of technologies that may be relevant for long-term data preservation and exploitation of data is constituted by semantic technologies. These are technologies that contribute in associating meaning to data and in more meaningfully organizing data, in meaningfully correlating data, as well as in converting data into information for more

effective decision making and in finding information that contextually relevant to users' needs [15]. They help with syntactic and representational as well as semantic interoperability. This general area of research is also getting renewed attention now that there is considerable excitement in the vision of the semantic web [16].

Central to the vision of the semantic technologies are ontologies. Ontologies are seen as facilitating knowledge sharing and re-use between agents, be they human or artificial [17]. They offer this capability by providing a consensual and formal conceptualization of a given domain. As such, the use of ontologies and supporting tools offer an opportunity to significantly improve knowledge management capabilities in large systems.

Ontologies are a means to complete the description of data, i.e. a technology able to provide a structured description of the digital object independently of the time (i.e. independent of the individual system technologies and information architectures): a digital object could be characterized by a digital ontology written using a standard relationship descriptive syntax. In time the ontology can be migrated onto new relationship encoding standards continuing in representing the structures within the digital entity. In this way also the structure of relationships among data is preserved.

A digital entity combined with the digital ontology can constitute the archival form on which long term preservation is based.

Examples@ESA

The following are examples of projects for which problems related to long-term data preservation and/or exploitation of Earth observation data are actual:

- **DILIGENT** – *A Digital Library Infrastructure on Grid Enabled Technology* – is an EC 6th FP project that will create an advanced test-bed allowing members of dynamic virtual organisations to access shared knowledge and collaborate in a, coordinated, secure, dynamic and cost-effective way. It'll be built by integrating both digital library and Grid technologies and will be able to serve different research as well as industrial applications. A test-bed is planned that will be demonstrated by two complementary real-life application scenarios of which one from the environmental e-Science domain led by ESA. The project starts September 2004 and lasts three years [6].
- **THE VOICE** – **THE**matic Vertical Organisation and Implementation of Collaborative Environments is an ESA General Studies Program study that analyses technologies for e-collaboration [19] and will demonstrate these by implementing several prototypes. Part of the study is the definition of a generic infrastructure for e-collaboration in the Earth Science domain. Technologies that are being analysed are Grid, Web-Services

and Semantic Web. Access to various types of data, their exchange and use is one of the major issues in the project. The project started April 2004 and will last 18 months.

- The “Grid On-Demand” portal developed at ESRIN gives NRT access to different level products of various sensors of given ESA Earth observation satellites. It defines a generic infrastructure where specific data handling and application services are seamlessly plugged in. Together with the high-performance data handling processing capability of the Grid, it provides the necessary flexibility for building an application virtual community with quick accessibility to data, computing resources and results. It integrates via Web Services access to the ESA catalogues and archive systems.

The Way Forward

We have discussed few issues related to technologies for long-term data preservation. The same technologies are fully relevant of enhancing the exploitation of the existing data holdings. A convergence of the two data utilisation views is more and more necessary for the proper handling of the huge volumes of environmental data available and planned.

For the Earth science community it is important to continue and invest in activities related to long-term data preservation. Projects like DILIGENT and initiatives like ERPANET / CODATA need full attention. It does not need to be said that non-accessibility and/or usability of data, missing information and/or knowledge is very expensive.

References:

- [1] Global Monitoring for Environment and Security Web-site: <http://www.gmes.info>
- [2] Oxygen, A new strategy for Earth Observation, *Earth Observation Quarterly*, June 2003, EOQ – N° 71, <http://esapub.esrin.esa.it/eq/eq71/chap1.pdf>
- [3] Envisat Web-site <http://envisat.esa.int>
- [4] Report on the NSF Workshop on Research Challenges in Digital Archiving: Towards a National Infrastructure for Long-Term Preservation of Digital Information, Workshop Report -- Draft 2.0 (Pre-Publication Draft), August 12, 2002
- [5] R. W. Moore, Preservation of Data, *SDSC Technical Report 2003-06*, San Diego Supercomputer Center, University of California, San Diego
- [6] C. Schwartz, Digital Libraries - Preservation & Digitization, Digital Libraries course, Simmons College, Boston, <http://web.simmons.edu/~schwartz/dl-preserve.html>
- [7] Final Report, The selection, appraisal and retention of digital science data, ERPANET / CODATA Workshop, Biblioteca Nacional, Lisbon, December 15-17, 2003
- [8] Global Grid Forum Web-site <http://www.gridforum.org/>
- [9] R. W. Moore, Preservation Environments, San Diego Supercomputer Center, University of California, San Diego

- [10] R. Suresh, K. McDonald, XML for Science Data Access, CEOS Joint Sub-Group Meeting, Frascati, Italy, 2002
- [11] C. Schwartz, Digital Libraries – LIS 462 Definition, Digital Libraries course, Simmons College, Boston, <http://web.simmons.edu/~schwartz/462-defs.html>
- [12] I. Foster, What is the Grid? A three-point checklist, *GRID Today*, July 22, 2002, <http://www.gridtoday.com/02/0722/100136.html>
- [13] European Data Grid Project Web-site <http://www.eu-datagrid.org>, EGEE project web-site <http://www.eu-egee.org>
- [14] Annex-I, Description of Work DILIGENT, Sixth Framework Programme of the EC, FP6-2003-IST-2, 30 April 2004
- [15] Semantic Association Identification and Knowledge Discovery for National Security Applications Web-site: <http://lsdis.cs.uga.edu/Projects/SAI/>
- [16] World Wide Web Consortium Web-site: <http://www.w3c.org>
- [17] Davies, I., Green, P., Milton, S. and Rosemann, M. (2003). Using meta models for the comparison of ontologies. In K. Siau et al. (Eds.), *Proceedings of the Evaluation of Modeling Methods in Systems Analysis and Design Workshop - EMMSAD'03*, Klagenfurt/ Velden, 16-17 June 2003
- [18] THE VOICE: THEmatic Vertical Organisation and Implementation of Collaborative Enviroments web-site: <http://www.esa-thevoice.org>

Information Mining in Remote Sensing Image – The KIM, KES and KIMV projects

Andrea Colapicchioni⁽¹⁾, Sergio D’Elia⁽²⁾, Mihai Datcu⁽³⁾

⁽¹⁾ *Advanced Computer Systems S.p.A. - Space Division*
Via della Bufalotta, 378
Rome (Italy)
Email: a.colapicchioni@acsys.it

⁽²⁾ *Earth Observation Applications Department, ESA Directorate of Earth Observation Programmes, ESRIN,*
Frascati (Italy)
Email: Sergio.DElia@esa.int

⁽³⁾ *Remote Sensing Technology Institute, German Aerospace Centre (DLR),*
Oberpfaffenhofen (Germany)
Email: Mihai.Datcu@dlr.de

THE PROBLEM: THE AMOUNT OF DATA AND THE NEEDED KNOWLEDGE

Modern imaging sensors, especially those aboard satellites, continuously deliver enormous amounts of data. The widespread of meter resolution images, is not only exploding the volumes of acquired data but also brings a new dimension in the image detail, thus fantastically growing the information content. These represent typical cases, where users need automated tools to discover, explore and understand the contents of large image databases. Without new theoretical concepts, methods and novel technological support, the existing volume of data prevents any systematic exploitation of Earth Observation (EO) data.

The state-of-the-art EO image catalogues allow only queries based on data like geographical position, acquisition date, sensor type, etc. and not on image content. Once the search parameters have been defined, these catalogues show the associated image to the user, leaving the interpretation task to his specific knowledge. Through his mental models and knowledge, the user interprets the image and derives meanings. Such a case configures a mono-directional information flow from the machine to the human being, preventing the opposite flow (user->system knowledge transfer).

There is a strong need to build up applications that help the user in image interpretation task, applications that permits to query the archives in content based mode, without having to know all the information contained in the images at signal level. They have to learn the conjectures made by the user during the interpretation task, and the concepts that are used in connection to the same image. This calls up a synergy between stochastic modelling, knowledge discovery, semantic representation, up to building a collaborative environment permitting to share the knowledge between heterogeneous user communities.

Image interpretation is not a simple task. Each user needs a set of accessory data, as for example GIS layers or text obtained through Internet. Yet, the amount of available information makes searches a demanding and expensive task. An environment where images are at the focal point, and where each user can navigate through a taxonomically structured knowledge, could be of extreme value.

The European Space Agency (ESA) has started since the year 2000 a series of research projects in the field, based on preliminary studies performed by DLR (German Aerospace Center) and ETHZ (Swiss Federal Institute of Technology, Zurich). This research activity is now bringing initial results on Content Based Image Selection and Information Discovery.

KIM: A FIRST STEP

Drawing upon the research experience of IMF-DLR, ETHZ and the systems engineering competence of Advanced Computer System SpA, an ESA Technology Research Programme (TRP) project was organised by ESA-ESRIN with the title: ‘Knowledge driven Information Mining in Remote-Sensing Image Archives’.

KIM represented a novel and unique theoretical concept and frame of collaborative methods for:

- Extraction and exploration of the content of image sets or other multidimensional signals
- Establishing the link between the user needs and knowledge and the information content of images
- Communicating at high semantic abstraction between heterogeneous sources of information and users with a very broad range of interest

KIM was the first prototype of a new generation of advanced tools and systems for:

- Intelligently and effectively accessing the information content in large EO data repositories

- Better exploration and understanding of Earth structures and processes
- Simplifying and enhancing access to and use of EO data.

The concept applied in KIM was aimed at building a system free from application specificity, so as to enable its open use in almost any scenario, and also to accommodate new scenarios required by the development of new sensor technology or growing user expertise. These goals were reached by defining a hierarchy of information representation levels, as depicted in the following figure, so as to enable the communication between the image archive and the users

KIM allows to:

- store the user’s specific knowledge;
- show to the same or other users not only the data, but also the semantics therein contained;
- mine the image archives using concepts;

This result has been achieved by extracting primitive image features (i.e.: texture, geometrical shapes, spectral information), and by providing the user with a simple, graphic interface to define weighted combinations of these features and to associate concepts (labels) to them through positive and negative examples.

Examples of semantically meaningful labels associated to concepts can be: “river”, “urban area”, “lacustrine vegetation”.

The weighted combinations of primitive image features and the associated semantic labels can then be applied to the entire dataset, and not only to the image from which they were defined.




Fig.1 – Information extraction

In 2003, other two projects were started by ESA-ESRIN:

- **KES** - EO domain-specific Knowledge Enabled Services
- **KIMV** - **KIM** Validation for EO archived data exploitation support

DOMAIN ONTOLOGY AND SEMANTIC MISUNDERSTANDING

Ontology is the specification of a conceptualisation. It can be related to a system (System Ontology, which can be domain independent and reused for different domains) or to a domain (Domain Ontology, specific for that domain).

Domain ontology is the set of definitions and concepts pertaining and belonging to a specific domain (and shared by concerned people).

Different domains have generally different ontologies. As an example, a climatologic expert could have a different vision of (and terms to describe) water compared to that of an oceanography expert.

In order to have an effective communication between different users, without risk of misunderstanding, they have to share a common ontology. A practical implementation of domain ontology can be through a semantic catalogue, which is fundamental for effective search and exploration.

Similar approaches (ontological/taxonomic) exist in e-commerce sites like Amazon.com and Yahoo shopping.

If it is known the domain of interest of the user who is going to define labels or semantic groupings, it is possible to derive a domain vocabulary, which defines the entities in semantic terms, and a classification of domains, which describes their inter-relationships.




Fig.2 – Links between objects

KNOWLEDGE TRANSFER: INTERACTIVE LEARNING

In the **KES** project we are refining the system that permits to associate semantic labels to weighted combinations of primitive image features. While defining labels, the user explicitly transfers knowledge to the system, which is stored and made available to the same or other users. In addition there is a further type of knowledge (implicit) that could be stored: if the user domain of interest is known, by observing the user interactions with the system during search and browsing, it is possible to infer the data of likely user interest and link it with the pertinence domain.

In a way similar to that of our memory, which builds up relations across synapses through experience, it will be possible to store in the system the user path during interactions.

Moreover, a relation could have a different weight depending on the number of times that a path is run with a positive feedback. Hence the weight represents the degree of correlations between system concepts.

Through these paths, and using their weights, it will be possible to predict the data of interest for a user that visualises one of these nodes. The weights allow defining a metric between ontology vocabularies, providing a method to compare learning path belonging to different ontologies. This peculiarity could permit to generate a system dictionary.

Hence we could define the set of a system acquired or learned knowledge as:

$$\text{Knowledge} = \text{Relations (classification)} + \text{Memory (Synapses)}$$

AGGREGATED LABELS: ACROSS SENSORS

In **KIM** it is simple to define a label by identifying a river through positive and negative examples. However, the “river” might become part of a wider concept (e.g.: water). This should be implemented without retraining the system for all possible water types. In **KES** a new kind of grouping, called “aggregated label”, will be introduced. In a similar way as in defining positive and negative examples on the image, it will be possible to define the concept “water” as:

Positive examples: sea + river + lake + water reservoir

Negative examples: streets + houses + mountains

This will create a new label, having the same searching and visualisation capabilities of any other label.

USED TECHNOLOGIES

The **KES** system is a distributed information system designed to provide an efficient, extensible catalogue system enabled with the acquisition of the user's knowledge to be used on network such as the Internet.

The **KES** system could be de-composed in two separate sub-systems:

- Ingestion Chain Modules: the chain that processes and stores the images into the Database Management System in order to extract the primitive features used in the definition of semantics from the images;
- **KES** server;

The **KES** ingestion chain is composed by several executables that, chained together, process each image, extracting the primitive features and inserting them into the database.

Because of the large number of images to be processed in **KES**, a Linux cluster approach has been chosen to allow to process all the images in a reasonable time.

The main **KES** server components could also be divided in three main tiers:

- Client: applications that run on the user's computer and provide the user and system administrator interfaces;
- Server: the functional modules that actually process the data. This middle tier runs on a server and provides data manipulation and exploitation tools;
- Database Management System (DBMS): stores the data required by the middle tier. This tier could run on a second server called the database server, or in the same middle tier server.

The **KES** system could be viewed as a web catalogue that provides standard functions, such as user management, image exploitation, user logging, with the addition of intelligent and specialised algorithms permitting to capture the user's knowledge and to define semantics from the data. All these functions are best viewed as application services, that is, user-driven specialised tasks performed by the system.

In order to minimise architectural design effort and risks, we need a set of instruments that provides an infrastructure to easily implement these services.

J2EE technology is the most effective currently available standard on the market to build applications, which are based on standardised, modular components, taking advantage of a complete set of services without complex programming.

With Java Web Start, the user launches applications simply by clicking on a Web page link. If the application is not present on his computer, Java Web Start automatically downloads all necessary files. It then caches the files on user's computer so the application is always ready to be re-launched at anytime -- either from an icon on the desktop or from the browser link.

The Database Server and Archive is based on the IBM/Informix Dynamic Server with the Geodetic Datablade Module, and is used to perform geo-temporal queries.

KIMV (KIM VALIDATION FOR EO ARCHIVED DATA EXPLOITATION SUPPORT)

The objective of the **KIMV** project is, starting from **KIM**, to implement and test a quasi-operational environment for simple access to enhanced image selection function (selection of images through a combination of standard spatio-temporal parameters and the information / features they contain).

This requires to:

- Identify the key primitive features and applications or services likely to benefit at most from enhanced image selection
- Identify and implement related **KIM** upgrade requirements
- Implement the interfaces with an existing archive, via disk storage or off-line media
- Ensure easy operability and adequate performance of the data ingestion chain in a quasi-operational environment
- Implement a generalised approach to envelope within the Service Support Environment (SSE) the available functions as services for simple user access to enhanced image selection
- Evaluate the performance of the system and its capability to fulfil user expectations in terms of image selection results (positive and negative)

KIMV provides an enhanced image selection service through SSE. The combination of standard parameters used in web catalogues (area of interest, time of interest, satellite, sensor etc.) and the **KIM** data mining capability enhance the possibility for the user to better define images that have to be processed from other **MASS** service providers.

The clusters and the class files are the key **KIM** data. They can be seen as a quantisation of the feature space of all the images in the archive and a coding of the archive information content. Thus being extremely sensitive for the quality of the mining process. Depending on the clustering algorithm in use, parallelisation of the data ingestion code and replication of data for efficiency may yield large benefits.

However, a global shared data structure, namely the cluster membership table, remains and must be managed centrally or periodically replicated and synchronised. The presence or absence of robust, efficient parallel clustering techniques determines the success or failure of cluster analysis in large-scale data mining applications in the future. Therefore two aspects have to be improved:

- **Throughput.** The "dyadic k-means" clustering algorithm used in **KIM** needs to be adapted to Linux Clusters in order to obtain the desired high throughput data ingestion.
- **Clustering.** A "living archive" would not benefit from the current clustering method (since new images might change the cluster "shape"). A different clustering approach, called "incremental clustering" is needed in such hypothesis in order to add new images to the archive without repeating the training for the already saved labels.

The major advantage of incremental clustering algorithms is that it was not necessary to store the entire data set in the memory. So, the memory and disk space requirements of incremental algorithms would be smaller than for traditional ones. Typically, they would be non-iterative. So their computation time requirements would also be smaller.

The motivation behind this work was that, in dynamic databases, items might be added and deleted over time. These changes should be reflected in the partition generated without significantly affecting the current clusters. Most of the incremental algorithms used before were order-dependent and their performance was proven only on relatively small datasets in the **KIM** case.

Thus, in the frame of this project new methods and algorithms have been developed to satisfy the needs of Image Mining technology, to be compatible with the **KIM** concept and architecture, and to ensure their information preservation.

RESULTS

The above-described projects are bringing initial results. They have demonstrated the possibility for a normal user to easily train the system via a graphic user interface on the searched feature and to identify all the images ingested into the system containing that feature (with a variable degree of probability). Of course, the types of detectable features depend on the ingestion methods, since it is possible to load the images at various scales or to extract a set of primitive features more than others. After the user has trained the system, the training results can be applied to the entire set of images also for other users.

This technique is been used, within the images ingested into the system, for:

- **Information Discovery**, which permits to obtain the identifiers of the images containing the feature for following order to or download from an archive.
- **Content Based Image Selection**, being tested on MERIS images, for the identification of the images free from clouds over the user area of interest. This permits to combine standard spatio-temporal (area, time, ...) and normal (mission, sensor, ...) parameters with other parameters related to the information contained in the ingested images. This has been implemented keeping track also of the localisation of the searched feature in the images, thus permitting to identify images having or not the selected feature in a specific region.

This technique is being tested in connection with an archive and related catalogue, on one side, and with the SSE, on the other. This organisation permits to create services from the archive, for example to identify the images containing the searched features, to order the images, to download the images, to obtain only the feature in GIS or map format, etc. Within the SSE it is easy to chain the services. Therefore it is possible to add further steps, for example to post process the data, combine it with additional non-EO data or information, to extract higher level information, etc.

CONCLUSIONS

The technologies for knowledge-driven image information mining are reaching a sufficient level of maturity for their integration into commercial products, as has been demonstrated here for a variety of remote-sensing applications. This opens new perspectives and offers huge potential for correlating the information extracted from remote-sensing images with the goals of specific applications.

These technologies shift the focus from data to information, meeting user needs, promoting scientific investigations, and supporting the growth of the value-adding industry, service providers and market, by permitting the provision of new services based on information and knowledge. They will also profoundly affect developments in fields like space exploration, industrial processes, exploitation of resources, media, etc.

The prototypes have demonstrated that:

- the results of advanced and very highly complex algorithms for feature extraction can be made available to a large and diverse user community

- the users, who can access the image information content based on their specific background knowledge, can interactively store the meta-information and knowledge

A new paradigm for the interaction with and exploitation of EO archives can be implemented, paving the way for much easier access to and much wider use of EO data and services

References

M. Datcu, H. Daschiel, et al, 2003, Information mining in Remote Sensing Image Archives: System Description, IEEE Trans. on Geoscience and Remote Sensing, ISSN 01 96 2892, Vol. 41, No. 12, pag. 2923 - 2936..

M. Datcu, K.Seidel, S. D Elia, P. G. Marchetti, 2002, Knowledge-driven Information- Mining in remote sensing image archives, ESA Bulletin 110, pp. 26-35. Available on <http://esapub.esrin.esa.it/bulletin/bullet110/dactu.pdf>

M. Datcu, M. Ciucu, Information theoretical approach for domain ontology exploration in large EO image archives, Proc. of IGARSS 2004.

The **KIM** prototype is available on <http://www.acsys.it:8080/kim>

SPASE – Space Physics Archive Search and Extract

Christopher Harvey⁽¹⁾, Jim Thieman⁽²⁾, Todd King⁽³⁾, D. Aaron Roberts⁽⁴⁾

*⁽¹⁾Centre d'Etude Spatiale des Rayonnements
9, rue du Colonel Roche, 31028 Toulouse CEDEX 4, France
Email: harvey@cesr.fr*

*⁽²⁾National Space Science Data Center
Code 633, NASA/GSFC, Greenbelt, MD 20771, USA
Email: james.r.thieman@nasa.gov*

*⁽³⁾Institute of Geophysics and Planetary Physics
UCLA, Los Angeles, CA 90095-1567, USA
Email: tking@igpp.ucla.edu*

*⁽⁴⁾Laboratory for Extraterrestrial Physics
Code 692, NASA/GSFC, Greenbelt, MD 20771, USA
Email: aaron.roberts@nasa.gov*

INTRODUCTION

Space Physics Archive Search and Extract (SPASE) is an international consortium of space physics data holding organisations which is developing a data model as well as recommendations for implementation of the data model. This will permit data centres to implement tools to allow science users to find space physics data of interest, intercompare the data, and retrieve selected data sets or portions of data sets. The search will be performed across multiple data centres through a single search query, and the results returned on a single results page. The effort involves creation of a common space physics metadata dictionary, and development of an intermediate level of software that will translate queries into the search mechanisms specific to each participating data centre ; and then return these query results in a common format. Thus, there is no need to change the software and search procedures in use at the individual data centres. Data set intercomparisons will facilitate determination of which data are useful for the purposes of the scientist-user and need to be requested from the data provider. Ultimately it should be possible to request and recover the selected data via this single user interface. The present status of the effort and plans for the future are discussed.

BACKGROUND

There is a growing awareness of the importance of preserving data gathered in connection with Space Plasma Physics research. Initially, data centres were set up to ensure long-term preservation of data for the support of future research. With the advent of the Internet, and more particularly the facility of communication offered by the World Wide Web, data centres can provide easy access to the archived data, and encourage its valorisation.

Access provision requires the procurement and maintenance of a suite of hardware, software and adequate network connections to enable each member of the international scientific community to find and recover those segments of the archived data required for his specific research. Valorisation consists of the provision of tools to facilitate and encourage the exploitation of this data with, in the first instance, the installation of Value-Added Services which allow the user to obtain the data in a standard form which he can readily exploit, as well as access to analysis tools which he may need.

The long term preservation of data poses new and special challenges :

- Firstly, it is necessary to preserve the data and the totality of the additional information required to make scientific use of this data : documents which describe the mission and the experiment, the syntactic (how it is written) and the semantic (what it represents) descriptions of the data, the auxiliary data sets (spacecraft orbit, attitude, etc.), other metadata, and possibly calibration parameters and necessary software for data usage.

- Secondly, it is necessary to solve the problem of being able to read over the long-term data which itself is recorded on supports which have only a limited life span. All data centres have addressed this problem, and there is now some convergence of view as to the best way of going about it. In particular, OAIS, “Open Archival Information System” [1], first conceived by the Coordinating Committee for Space Data Standards (CCSDS), is now the ISO (International Standards Organisation) standard 14721.2. This conceptual model provides a systematic analysis of the problems of long-term archiving.

Long-term archiving of the data and everything which is necessary to exploit it permits, in principle, the study of :

- the evolution of Sun-Earth relations,
- the compatibility of new theories with existing data,
- the application of new analysis tools to existing data.

Furthermore, access to multiple data sets allows

- easy comparison of the same or related phenomena as observed by different instruments.

This last point is the essence of SPASE. For reasons developed in the next section, Space Physics data is not, and never can be, archived at a single site. The purpose of SPASE is to make the end users' local data system appear as a single coherent global system, while in reality it is part of a system consisting of many geographically dispersed centres, storing data from different missions or different experiments, possibly in different formats, and offering different value-added services.

Different scientific disciplines have their own terminology and hierarchy of concepts which will be used to describe data and services. It is therefore important to define what is meant by Space Plasma Physics. SPASE considers this to be physics of regions in the heliosphere from the corona to the interstellar shock, including the ionised environments of the Earth and other planets, their satellites, and other bodies of the solar system. Nevertheless, one of the important requirements of SPASE is that technological developments be, in principle, applicable to a wider range of disciplines.

THE GLOBAL DATA SET

There are several reasons why no one data centre can archive all data :

1. “All data” from what discipline ? However disciplines are defined, there will always be borderline cases.
2. Even when the disciplines are defined to the best of our ability, it rapidly becomes clear that the quantity of data to be archived far exceeds the ability of any one centre to handle it.
3. Likewise, no one agency will wish to take responsibility for funding one massive centralised data centre.
4. Data centres do more than just store old data : they update the documentation in the light of user requests, and they provide services to the users. The development and evolution of these services, essential for the valorisation and exploitation of the archived data, require the participation of an active local (in a nearby time zone and speaking the same language) scientific research community, preferably the community responsible for producing the data in the first place.
5. The preceding point is also applicable to services which are not related to any particular archived data set : these include such services as the maintenance and provision of software libraries, the maintenance of services related to orbitography and the identification of satellite rendezvous. These services are extremely important, and require significant investment of both scientific and technical effort. They must necessarily be exploited nearby to the community responsible for their development.

For all these reasons data centres will become increasingly dispersed across the globe.

The end user wants information, data, documentation, and services, but is not really too worried about where these are to be found, so long as they correspond to his needs and are of good quality. This is what interoperability is all about. The ultimate objective is that the data system appears to the end user as a single coherent global system, while in reality it consists of many geographically dispersed centres, storing data from different missions or different experiments, possibly in different formats, and offering different value-added services. This goal is termed “Data Centre Interoperability” ; it is also known (especially to astronomers) as the “Virtual Observatory”.

Today we are a long way from this goal. There are a considerable number of data servers and the user, to make a thorough search for data which may be useful for him, must connect to each individual server, browse the pages offered using an interface and an interrogation procedure both of which are different on each server, and eventually request from each individual centre the data he wants.

THE ORIGIN OF SPASE

The need for interoperability in space physics was clearly identified during an ISTP (International Solar Terrestrial Physics) meeting in September 1998. As a result of that meeting, a small group of the world's largest space physics data centres worked together, and in 2002 they signed a charter to work together in a collaboration now called "Space Physics Archive Search and Extract". (The original title "Space Physics Archive Search and Exchange" has recently been changed so as to emphasise that SPASE will allow data retrieval by any user). The charter, available from the consortium Web site [2], states the objectives of SPASE :

1. Build a system that identifies and meets the global Space Physics User Requirements.
2. Facilitate the circulation of Space Physics scientific and technical information.
3. Facilitate the (two-way) interface with international organisations responsible for technical standards and developments.
4. Avoid duplication of effort between geographically distant space physics archiving centres.
5. Ensure the compatibility of the architectures used for the global distributed system and as many of the individual data centres as possible.
6. Minimise costs by development of a system that :
 - uses widely approved technical standards,
 - is easy to maintain (both globally, and in each archiving centre),
 - can easily evolve so as to profit from likely future technological developments,
 - can be interfaced with the systems of adjacent disciplines and, hopefully one day, with a global system.

The objective of SPASE is to demonstrate interoperability and develop the tools that are essential to be able to realise it for the Space Plasma Physics community. It is not the objective of SPASE to run any data centre ; this remains the responsibility of each of the separate data centres.

When fully developed, "Interoperability" must allow the user, wherever he is situated, to find via a graphical interface with which he is familiar :

- the information he wants,
- the data he needs, and
- the tools to exploit it,
- all in a form which he can readily use.

SERVICES OFFERED

The objective of data centre interoperability is to provide the end user with homogeneous access via his local data centre to the information, data, and services which are today provided separately by one or more centres, each using its own specific user interface. Every data centre offers a variety of data and services, the richness of the offer depending upon the size of the centre and the resources available ; although sometime a small centre may offer a very specialised service.

It is useful to present briefly the services currently on offer, or under development, and which are possibly to be included in SPASE interoperability.

Data Searching

The first step towards making a data request is to find the data to be requested. Individual data centres provide many ways of navigating to find data (e.g., often within the mission/experiment/dataset hierarchy). SPASE will offer the possibility of searching by keywords and, of course, time interval but may not have as extensive or specific a hierarchy of keywords or metadata as some data centres use. The results obtained by automatic interrogation of the remote data centres are presented to the end user in a homogeneous way, that is, the results are, as much as possible, mapped into a common terminology. To optimise the performance, some considerable effort has been spent in drafting the SPASE Metadata Dictionary.

It is realised that not all data centres are capable of searching and replying to the questions posed with the same level of detail : account is taken of this in the presentation of the results, so as not to overload the page with “hits” of little relevance to the question originally posed by the user.

Data Request and Recovery

Most space physics researchers are interested in obtaining data from time intervals of particular interest for their current research. Once a potentially interesting dataset has been identified, the user will want to recover that part of it spanning his time interval of interest. This is the next priority for SPASE ; it should be relatively easy to implement if the data is delivered (to the end user) in the files in which it is archived : that is, the data files are not opened, there is no time segment extraction and/or file concatenation.

Value-Added Services

It is the Value-Added Services which encourage the scientific exploitation of the archived data and thus distinguish a data centre from a simple archiving centre. This is a vast exercise which must perpetually adapt itself to satisfy the evolving requirements and aspirations of the scientific user community. Whilst some scientists wish only to recover data in a convenient format and with sufficient documentation to be able to use their own data analysis tools, others are very interested in value-added services.

Amongst the services which may be provided are :

- The selection, from within the archived data files, of data spanning the precise time interval requested by the user.
- The extraction of particular parameters requested by the user from the multitude of parameters contained within certain data sets.
- Conversion of the selected data to a format adapted to his own local analysis system (CDF, netCDF, ASCII, *etc.*). This highly desirable functionality may depend upon the capability of the centre holding the data to perform the transformation, but SPASE may provide some format translation capability as well.
- The search for data in association with external criteria, such as the direction of the interplanetary magnetic field, or the index of geomagnetic activity.
- Transformation of the coordinate system of vector and tensor quantities.
- User-defined graphical displays. The ability to visualise and compare data from different missions would be especially useful.
- Documentation and bibliographic reference services.
- Software libraries, from which users can download code (plus installation and test procedures, and full documentation) so as to have, in their own laboratories, well-validated implementations of standard analysis procedures.
- Access to the Satellite Situation Center. This centre, operated jointly by the NASA/GSFC Space Physics Data Facility (SPDF) and the National Space Science Data Center (NSSDC), provides a tools to identify satellite rendezvous, times at which spacecraft are on the same magnetic field as some ground-based observational facility or even other spacecraft.

CURRENT SPASE ACTIVITY

SPASE Data Model

A data model is a set of terms and their relationships that capture the essential concepts of a given domain. The SPASE data model is intended to give a quite complete description of the domain of solar and space physics products resulting from observation and modelling. The products include datasets, images, software, and documentation. This data model includes terms relevant to all “Sun-Earth Connection” domains, but it does not include Earth Science terms. The set of terms and their definitions are referred to as the Metadata Dictionary.

More specifically, the data model should :

- Provide a way of registering products using a standard set of terms that allows the products to be found with simple searches.
- Allow searching for products containing particular physical quantities (*e.g.*, magnetic field, spectral irradiance) that are represented in a diverse array of data products.
- Create a means of mapping comparable variables from many products onto a common set of terms so that visualisation and analysis tools can be used on all of them without regard to the origin of the data.

The accomplishment of these tasks requires “middleware” that understands product registries and performs the translations needed to map the idiosyncratic file names and quantities of each repository onto standard terms. This intermediate layer, which can take many forms for different purposes, will provide the links necessary to connect user applications and search-and-retrieval front ends to data repositories. Ultimately, the data environment centred on the data model will involve a number of software tools as well, linked together as internet-based services or other means. Specific software tools and documentation associated with products will be accessible straightforwardly. This “system” has the potential to provide capabilities that can aid even expert users of a particular dataset (*e.g.*, on-the-fly coordinate transformations ; the ability to merge datasets from different instruments), in addition to providing the broad access needed to solve global problems in Sun-Earth connection physics. Success will require a concerted cooperative effort across disciplines. Existing efforts in Space and Solar Physics as well as in other areas such as Earth Science and Astronomy will guide the work.

Metadata Dictionary

The metadata is data which describes the data. This is used for two distinct purposes :

- to find data within archives, and
- to exploit data once it is found and recovered.

The requirements are not the same for these two purposes, although there is clearly considerable overlap. The metadata should be complete, that is, contain all the information required to read the data, to understand the data (syntactic description), and to understand what it represents (semantic description), including how it was obtained ; the latter impacts upon the scientific significance of the data.

Metadata should be held in a homogeneous way, and the various data centre applications, such as search, data extract and preparation for delivery, interface to data centre applications, *etc.*, should select the parameters which they need. A data centre search engine, as well as the search possibility offered to external data centres (for example, via Web services), thus has, in principle, the possibility of searching any or all of the metadata. This does not imply that it is necessary to implement searching of all the information held ; merely that this possibility exists and may be implemented if one day it becomes desirable. Nevertheless, searching through the totality of the metadata will be relatively easy if the metadata is preserved in a well-ordered way which allows use of generic software.

The metadata dictionary is clearly specific to any particular discipline, and it will probably evolve rather slowly.

Technical Implementation

Any software which is developed specifically for some discipline, such as Space Plasma Physics, will rapidly fall into disuse unless it is maintained. Furthermore, maintenance will not be negligible because the technology is evolving rapidly ; and so too are the user expectations. For both these reasons there has been less SPASE activity so far on the technological front, especially considering that SPASE consortium participants were not, until just recently, receiving any funding for direct support of SPASE.

But some progress has been made. For example, a breadboard model has been developed which allows the user’s local centre to prepare a request in a standard “HTML-type” format, which is then sent to other centres. These centres can understand the question, search their data holdings using their own specific search tools, and send their reply in the form of an XML page. The user’s server then collates the information received in the different XML replies, and produces a single homogeneous page presenting the results of the search. The initial effort was modest (only two data centres involved) compared to the requirements, but this breadboard has highlighted a number of problems which must be resolved.

Each data centre must be able to answer the questions received. This seems to indicate that a structured digital description must be associated with each archived data set, which is currently not true in all data centres. The situation will undoubtedly improve in the larger centres, but many new, smaller, data centres will appear in the coming years. The assumption that if a data centre cannot respond to a query then everything (rather than nothing) is selected leads each locally inapplicable search criterion to supply more “hits” than a well-targeted question which is handled completely. When compiling the “results” from the received XML replies, account must be taken of the fact that less discriminating centres provide proportionally more hits, most of which are not useful.

Most of the data centres which will first be connected using SPASE are “legacy” data centres ; that is, they were built before SPASE, and they have their own keywords. A tool is needed to translate between the keywords used in different data centres. This is more than just a translation table : a framework is needed into which the different words can be fitted, in order to permit “sloppy translation” when exact translation is impossible between two lists of keywords.

Finally, due to the enhanced network security, the breadboard which worked at the start of 2001 has had to be adapted to today's security requirements.

Difficulties Encountered

The principal difficulties encountered are related to the very high expectations of the community of scientific users, coupled with the large amount of preparation required to develop a system which is sufficiently robust to fulfil these expectations. This preparation involves a very wide community, involving both scientists to define the data dictionary, and engineers to develop tools to exploit that dictionary, and deal with other problems such as resolution of conflicts (arising, for example, from receipt of contradictory information from two different sites) which are inevitable within all distributed network information systems.

Interoperability, standards for data and for metadata, protocols for communication and exchange of information between data centres, all these things need to be discussed openly. World-wide scientific organisations such as COSPAR, IUGG, do occasionally organise meetings where these issues are presented, but they do not provide a forum for work to be done. There are at least two technical bodies supported by the International Council of Scientific Unions which are specifically mandated to study these subjects :

- CODATA (Committee on Data for Science and Technology, <http://www.codata.org/welcome.html>), and
- ICSTI (International Council for Scientific and Technical Information, <http://www.icsti.org/>).

These international organisations are too general to be concerned with the particular requirements of Space Plasma Physics. On the other hand, technical solutions, and especially any tools used by Space Plasma Physics, must be used by other disciplines : otherwise, maintenance costs will ensure that their lifetime be short. Scientists need to work with scientists and engineers of other disciplines,. All these activities, international scientific coordination, international technical coordination and coordination of the scientists and engineers at all levels, must be pursued in parallel ; while all the time the scientific users are impatient to have a solution which satisfies their expectations as soon as possible. The temptation for individual groups of scientists to “go it alone” is very great but, in the long term, globally counter-productive.

Data archive interoperability also has a problem of a completely different nature : resources. Developments in information technology are certainly not research in plasma geophysics. Are they research in information systems technology ? Interoperability needs tools which are stable in time, and those at the leading edge of technological development may not be in this category. Both disciplines, geophysical research and information systems research, find that support for this activity should come from “elsewhere” ! Furthermore, some space agencies tend to be “Project oriented”, which make it difficult to develop an archive which is multi-mission, even though this is the only tenable solution for a long-term archive. The solicitation of support outside the space science community generally often meets with the objection that it is a disguised request for science support. The diversity of the funding, international, national and regional, for the existing data centres does not simplify international cooperation.

DISCUSSION AND CONCLUSION

Interoperability with respect to data search and extraction is just beginning, and the Value-Added Services listed above are far from being fully implemented. Priority should be given :

1. To services other than simple data location : data selection, data recovery, and possibly transformation before delivery. The obvious transformations which should be offered are
 - Data resampling, which is the application of “standard” algorithms to change the frequency at which data is sampled ; this includes, for example, synchronisation of data sets to facilitate comparative analysis, and averaging to reduce the volume of data to be shipped.
 - Transformation of coordinates, for example to convert vector quantities such as the magnetic field from the coordinate system in which is archived to some other coordinate system more suitable for the subsequent analysis.
 - Transformation of data format prior to delivery, for example, to one of the (too many) “standard” formats used by the different analysis tools of the discipline.
2. To compatibility with other disciplines : verification that the data model developed for SPASE can be applied elsewhere, first to neighbouring disciplines, then further afield. This is essential for the development of the “Virtual Observatory” concept, and also for the long-term viability of tools for metadata handling and interoperability.

In conclusion, the problem of the access to and exchange of information is one of today's major technological challenges. Storage of digital information is becoming cheaper every day, but there is little point in storing data if it cannot be found by the people who need it, when they need it. This is not a problem specific to space plasma physics. Space Plasma Physics has identified its particular problems, and SPASE is a space physics initiative for solving these problems. It is certainly not the only such initiative : for example, the International Virtual Observatory Alliance [5] is another. What is really needed is interaction with neighbouring scientific disciplines which surely have similar problems, and with the groups responsible for the development of technology and of standards, so as to ensure that progress is made in an as efficient and durable a manner as possible.

REFERENCES

- [1] Consultative Committee for Space Data Standards, “*Reference Model for an Open Archival Information System*”, (OAIS). Blue Book, Issue 1, January 2002. ISO standard 14721. Available from <http://www.ccsds.org/documents/650x0b1.pdf>
- [2] The Space Physics archive Search and Extract (SPASE) site has the home page <http://www.igpp.ucla.edu/spase/>
- [3] C.C. Harvey, C. Huc and M. Nonon-Latapie, “Centre de Données de la Physique des Plasmas”, *Adv. Space Res.* Vol. 31, pp. 1291-1295, 2003. Available from <http://www.sciencedirect.com/> by selecting journal and the volume.
- [4] C.C. Harvey and C. Huc, Future Trends In Data Archiving and Analysis Tools, *Adv. Space Res.*, Vol. 33, pp. 816-821, 2004. Available from <http://www.sciencedirect.com/>
- [5] The International Virtual Observatory Alliance (IVOA) home page : <http://www.ivoa.net/>

