INTELLIGENT ARCHIVE CONCEPTS FOR THE FUTURE

H. K. Ramapriyan a, *, G. R. McConaughy a, C. S. Lynnes a, R. Harberts b, L. Roelofs b, S. J. Kempler a, K. R. McDonald

a NASA Goddard Space Flight Center, Greenbelt, MD 20771
(Ramapriyan, Gail.R.McConaughy, Christopher.S.Lynnes, Steven.J.Kempler, Kenneth.R.McDonald)@gsfc.nasa.gov
b Global Science & Technology, Inc., 6411 Ivy Lane, Suite 300, Greenbelt, MD 20770
(harberts, roelofs)@gst.com

KEY WORDS: Archiving, Artificial Intelligence, Remote Sensing, Global-Environmental-Databases, Automation, Knowledge Base, Decision Support, Data Mining

ABSTRACT:

Sponsored by NASA's Intelligent Systems Project, a conceptual architecture study is under way to address the problem of getting the most societal value from the large volumes of scientific data that NASA expects to accumulate in the future. Beyond improvements in hardware technologies, advances are needed in concepts and tools to enable intelligent data understanding and utilization. Some of the challenges besides large and ever-growing volumes of data are: data acquisition and accumulation rates tend to outpace the ability to access and analyze them; the variety of data implies a heterogeneous and distributed set of data providers and users; unassisted human-based manipulation of vast quantities of archived data is intellectually overwhelming and cost prohibitive; for applying NASA technologies to operational agencies' decision support systems, it is necessary to demonstrate feasibility of near-real-time utilization of vast quantities of data and the derived information and knowledge; and future data access and usage are difficult to anticipate. The objective of the study is to formulate ideas and concepts and to provide recommendations that lead to research by the computer science community in the near-term, prototyping to demonstrate feasibility in the mid-term, and operational implementation in the period from 2012 to 2025. An abstracted architecture is defined for an intelligent archive showing functionality without regard to physical distribution. The architecture shows significantly enhanced functionality anticipated by NASA as required to serve its and society's future needs. This expression of functionality can help target research by the computer science and information technology communities.

1. INTRODUCTION

One of NASA's vision statements is "To improve life here on Earth." Derived from this are several strategic objectives for the Earth Science Enterprise (ESE). The ESE supports NASA's vision by providing, to the research community, policy makers and the general public, data and information products and the knowledge derived therefrom to enable/improve decision-making. Some of the benefits of this are: conserving resources, increasing prosperity, improving the quality of life, reducing impacts of disasters and saving lives. To enable this, ESE combines the NASA-unique capabilities for space-borne observations with research in various Earth science disciplines to address a number of scientific questions (NASA, 2000b). A suite of Earth observing satellites, funded research investigations, and the Earth Observing System Data and Information System (EOSDIS) along with the federation of Earth Science Information Partners (ESIPs), including the Distributed Active Archive Centers (DAACs), are now operating to provide an unprecedented amount of data and information products to a broad user community. In addition, ESE supports research in and development of technologies needed for future Earth observing systems (NASA, 2002a), research in and development of information systems' technologies (NASA, 2002b) and an application program (NASA, 2002c).

As expressed in NASA's plans referenced above, some of the key science and applications goals for 2010 are to improve predictive capabilities for:

- Weather
  - 5-day forecasts with over 90 % accuracy
  - 7-10 day forecasts with 75% accuracy
  - Routine 3-day forecasts of rainfall
  - Hurricane landfall prediction with +/- 100 km accuracy 2-3 days ahead of time
  - 2-day forecasts of air quality
- Climate
  - Routine 6-12 month seasonal predictions
  - Experimental 12-24 month predictions
  - 10-year experimental climate forecasts
- Natural Hazards
  - Continuous monitoring of surface deformation in vulnerable regions with millimeter accuracy
  - Improvements in earthquake and volcanic eruption forecasts
  - Improved post-eruption hazard assessment

Some of the areas where additional progress is speculated as a vision for the future (2025) are (NASA, 2000a):

- 10-year climate forecasts
- 15- to 20-month El Niño prediction
- 12-month regional rain rate
- 60-day volcano warning
- 10- to 14-day weather forecast
- 7-day air quality notification
- 5-day hurricane track prediction to +/- 30 km
- 30-minute tornado warning
- 1- to 5-year earthquake experimental forecast

More details about current thinking, identification of research and measurements needed, and projections on what could be achieved are given for three key areas of significant societal impact: biological invasion and ecological forecasting (Schnase et al, 2002), understanding sea level changes (Chao et al, 2002), and understanding and responding to earthquake hazards (Raymond et al, 2002). For the society to derive benefits from these advances, not only is it necessary to conduct the research

* Corresponding author.

and development to enable such capabilities, but it is also essential to demonstrate such technologies in operational environments and transition them to operational agencies. In most cases, the operational agencies have requirements to use the information and knowledge acquired through such advancements in real- or near-real-time environments to support decisions that have significant impact on society. Thus, for the scientific and technological advances from NASA to be truly applied for the benefit of the community, it requires advances in sensor technologies, satellite systems, scientific research, transformation of data to information and to knowledge, and dissemination of knowledge to support decisions in a timely manner.

Over the past decade, there have been significant advances in our ability to collect, archive and disseminate data. In the 1980s NASA's Earth science data were generally held by principal investigators or held at specialized data systems with virtually no cross-interaction or interoperability. It was difficult to find data unless a user knew where the data were. The access to data became easier, and the quality of services associated with data increased significantly with the development of Version 0 EOSDIS (Ramapriyan and McConaughy, 1991) through interoperable, geographically distributed, data centers. The capabilities of the World Wide Web (WWW) and its widespread use have revolutionized the access to information. Considerable progress has been made in ingesting, archiving, and distributing large volumes of data using distributed databases with EOSDIS (Ramapriyan, 2002; Moore and Lowe, 2001; Moore and Lowe, 2002), discovery of the existence of datasets and services through the Global Change Master Directory (Olsen, 2000; Smith and Northcutt, 2000), the area of interoperability through the EOS Data Gateway (EDG) (Pfister, 2001), Data and Information Access Link (DIAL) (McDonald et al, 2001), Distributed Oceanographic Data System (DODS) (Cornillon, 2000), Alexandria Digital Library (Frew et al, 1999), EOSDIS Clearing House (ECHO) (Pfister, 2001), and other efforts. Access to specialized data products and applications' development for focused user communities has been enabled by NASA through the Federation Experiment involving over 24 Earth Science Information Partners (ESIPs) (NASA, 1999; ESIPFED, 2002). There are a number of efforts underway to take advantage of distributed computing and storage resources that are generally referred to as "Grid Architectures." Examples of these are: National Science Foundation's National Technology Grid (Smarr, 1998), NASA's Information Power Grid (Thigpen, 2002), U.S. Department of Energy's DISCOM (DISCOM, 2002), GriPhyN (Avery and Foster, 2002), NEESgrid (NEESgrid, 2001), and the Particle Physics Data Grid (PPDG, 2002).

Even with the above accomplishments, further developments in information sciences and technology are critical to the achievement of NASA's vision for the future and its deployment for operational applications. Currently, NASA's Earth Science Technology Office supports the development of information system technologies (NASA, 2002b) for near- to medium-term deployment in mission applications. In addition, development of more embryonic technologies for longer-term mission adaptation is supported through other NASA programs, addressing needs of all NASA enterprises. An example of such a program is the Computing, Information, and Communication Technologies (CICT) Program. A component of the CICT program is the Intelligent Systems Project (ISP). (See NASA, 2002d). One of the technical areas under the ISP is Intelligent Data Understanding (IDU). Several basic research activities and conceptual studies are being supported within IDU. This paper reports on one of the conceptual studies within IDU, namely, Intelligent Archives (IA).

The motivating factors for this study are:

- Data acquisition and accumulation rates tend to outpace the ability to access and analyze them. For example, the rate at which EOS data are accumulating in the archives today is about 3 TB/day.
- Beyond the obvious needs for more efficient storage and access to data that are met by improvements in hardware technologies, advances are needed in concepts and tools to facilitate intelligent data understanding and utilization.
- The variety of data implies a heterogeneous and distributed set of data providers that serve a diverse, distributed community of users.
- Unassisted human -based manipulation of vast quantities of archived data for discovery purposes is intellectually overwhelming and certainly cost prohibitive.
- While there is no substitute for human intelligence, it is necessary to provide automated "intelligent assistants" to maximize the utility and application of data and transforming them into information and knowledge.
- Especially if NASA is to migrate its technologies to operational agencies' decision support systems, it is necessary to demonstrate the feasibility of near-real-time utilization of vast quantities of data and the derived information and knowledge.
- The types of data access and usage in future years are difficult to anticipate and will vary depending on the particular research or application environment, its supporting data sources, and its heritage system infrastructure.

The past and present advances in data management mentioned above and the expected advances in hardware and software/information systems technologies over the next decade provide exciting possibilities for improved access and utilization of data as well as design and deployment of true "knowledge building systems (KBSs)," moving beyond the basic capabilities of "Data and Information Systems." A KBS in the context of Earth sciences can be viewed as an end-to-end system starting with sensors (space-borne, airborne or Earth-bound) and ending with users who derive knowledge through scientific research and/or exploit the knowledge in real-life applications. The knowledge itself is preserved for posterity. Knowledge building involves a dynamic interplay between people and technology that transforms observations into data initially. The dynamic interplay then transforms data into information and information into knowledge. An IA fits within this end-to-end context and will support this knowledge building enterprise with new capabilities and facilities.

The purpose of this paper is to introduce the IA and show its architecture at a conceptual level. The next section covers the definition of a few basic terms and a discussion of the architecture. Section 3 provides an operational concept for an IA. Section 4 concludes the paper with a summary and recommendations for future work.

2. CONCEPTUAL ARCHITECTURE

The term IA needs some explanation. The term archive means permanent data storage. However, it is common practice to include functions in support of ingest processing, indexing and cataloging, quality assessment, product generation, search and retrieval, and delivery of selected data. In this paper, we restrict the domain of the IA to storage but include all items stored to support “end-to-end” research and applications scenarios. The physical locations of stored items are irrelevant to this discussion. They could be anywhere in the end-to-end system, including on-board caches with the sensors and those in the client systems with the users. Stored items include:

- Data, information, and knowledge
- Software needed to manage holdings
- Interfaces to algorithms and physical resources to support acquisition of data and their transformation into information and knowledge, storing the protocols to interact with other facilities

Thus, the phrase IA generalizes the term archive from a simple repository of data to one that supports and facilitates derivation of information and knowledge. Data, information and knowledge are defined below as in (Ramapriyan et al, 2002a):

- Data: an assemblage of measurements and observations, particularly from sensors or instruments, with little or no interpretation applied. (e.g., measurements from scientific instruments, market’s past performance)
- Information: a summarization, abstraction or transformation of data into a more readily interpretable form. (e.g., results after performing transformations by data mining, segmentation, classification, etc., such as a Landsat image spatially indexed based on content , assigned a “class” value and subset for an application, or National Weather Service storm monitoring fused with a map of the spatial location of the Washington D.C. Beltway)
- Knowledge: a summarization, abstraction or transformation of information that increases our understanding of the physical world. (e.g., predictions from forward runs of models, published papers, output of heuristics, or other techniques applied to information to answer a “what if” question such as “What will the accident rate be if an ice storm hits the Washington D.C. Beltway between Chevy Chase and the Potomac crossing at 7 a.m.?”)

The word intelligent in IA implies certain characteristics that increase the ability of an archive to operate more autonomously than a “dumb” archive and provide better service to users (as an intelligent assistant) with less operator intervention and, hence, lower cost. Some of these characteristics are shown below, grouped under three different categories of autonomy:

- Holdings’ Management Autonomy
  - Provides data to a science knowledge base in the context of research activities
  - Is able to exploit and use collected data in the context of a science enterprise
  - Is aware of its data and knowledge holdings and is constantly searching new and existing data for unidentified objects, features or processes
  - Facilitates derivation of information and knowledge using algorithms for Intelligent Data Understanding
  - Works autonomously to identify and characterize objects and events, thus enriching the collections of data, information and knowledge
- User Services Autonomy
  - Recognizes the value of its results, indexes/formats them properly, and delivers them to concerned individuals
  - Interacts with users in human language and visual imagery that can be easily understood by both people and machines
- System Management Autonomy
  - Works with other autonomous information system functions to support research
  - Manages its activities and functions from sensor to user
  - Manages the optimization of its own configuration
  - Observes its own operation and improves its own performance
  - Has awareness of the “state” of its cooperating external partners

Figure 1 shows the end-to-end context for an intelligent archive. Note that the end-to-end system is highly distributed. In fact, all but the smallest intelligent archives will be built upon a geographically widely dispersed set of physical resources and information repositories. Thus, our conceptual architecture will be implemented over a “grid” of computing resources: IA can be viewed as a level of software applications and logical organization applied to the grid. The presence of feedback loops connecting intelligent algorithms, methods, and technologies within the archive is an important feature. When implemented throughout an enterprise system, these automatic feedback loops create adaptive systems that respond to the need for greater data utilization and improved digital scientific services.

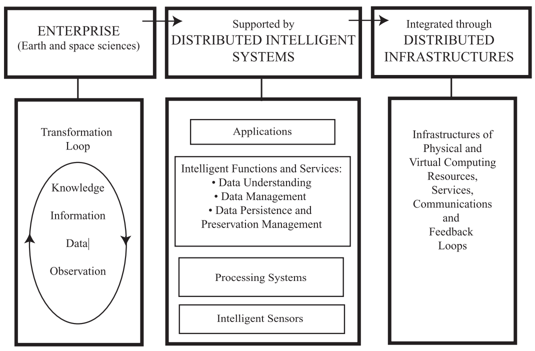

Figure 1. Context for an Intelligent Archive

The functions of the Intelligent Archive can be deployed over discrete computing facilities or optionally over a network of facilities as an open, distributed resource. Small to large intelligent archives in the near future can be built and operated either collaboratively or privately because of evolving technologies, standards, and methods. Rapidly changing computing technology will permit the formation of new enterprise system capabilities built upon distributed interconnected infrastructures - a system of systems. With the inclusion of more self-directed automated processes for autonomous discovery, self-aware management, and self-healing, new advantages can be realized for both archive users and archive operators.

A functional view of the IA architecture is shown in Figure 2. This figure shows the core functional capabilities of an IA. Here, the only distinction between permanent and interim archives is that permanent archives preserve scientific data, information, and knowledge indefinitely while interim archives establish policies on the duration of retention. Both interim and permanent archive systems draw upon at least five core functions. The rounded rectangular boxes represent core system functions with corresponding sub-functions listed in dashed boxes. Intelligent systems and data understanding technology underlie many of the core archive functions. Future operations’ functions will be characterized by intelligent capabilities that

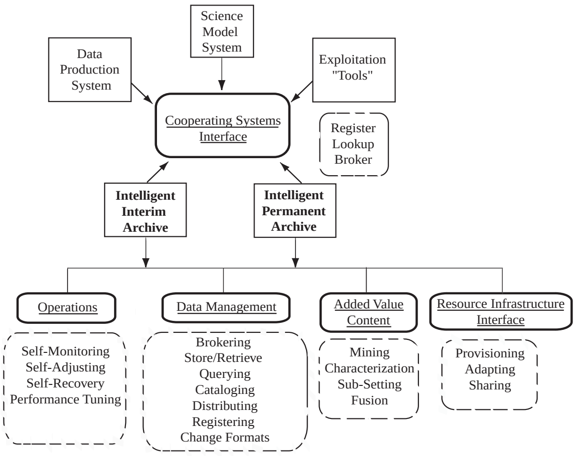

Figure 2. A Functional View of an Intelligent Archive

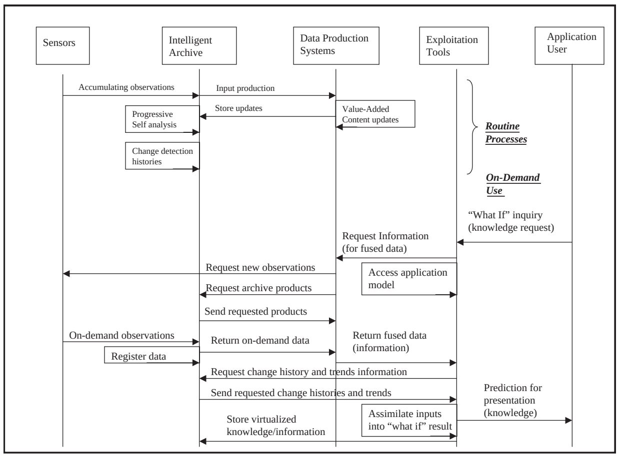

Figure 3. Abstracted Operational Concept for an IA

increase automated self-regulation, adjustment, recovery, and performance tuning. Intelligent data management consists of such functions as semantic queries, retrieval strategies, registering and cataloging assets for human and machine level access, and learning access patterns for improving storage/retrieval strategies.

3. OPERATIONS CONCEPT

Scenario-based perspectives can stimulate thinking about the capabilities needed in an IA. With this point of view, two Earth science scenarios (precision agriculture and advanced weather prediction) have been explored in this study to uncover requirements for intelligent system services and capabilities. Details of these are beyond the scope of this paper (see Ramapriyan et al, 2002b). An abstracted operational concept based on these practical scenarios is shown in Figure 3. This figure displays an enterprise-wide flow of information horizontally among the high-level system components that transforms observations into knowledge. The figure follows conventions for event diagrams using descending vertical lines to indicate time increasing. Interactions below each object show how the component objects collaborate for two illustrative aspects of the operational scenario: routine processing and on-demand use. There are many other interactions supported by an IA to cover the various functions discussed in section 2, but they are left out in this figure for purposes of clarity and readability.

Observations accumulate in an archive, which then provides data to production systems. These will, in turn, store their products in an archive. This operational flow is managed automatically through intelligent functions for coordinated interoperable processing. Moreover, an intelligent archive also automatically analyzes its contents to build and update archive metadata. For that operation, it uses data characterization algorithms to detect changes and update the archive's metadata.

Application users have flexible, on-demand access to data and information concurrently with routine operations. For example, suppose a scientist is developing a high-resolution model of a weather event and needs to perform a "what if" inquiry to test predictive accuracy of the model. The scientist uses an application tool interfaced with the knowledge building system to formulate an inquiry. The application tool can then select the required services and issue requests against real or virtual archive holdings. Requests for virtual holdings trigger a cascade of requests to obtain, for example, fused data, or subsets of observations, or the output from another predictive model or even new observations. Though not shown in the figure, that cascade of actions across these cooperating nodes would require the IA to be a repository of protocols of the permissible rules for collaboration, and might even require storing an image of the dynamic "state" of its collaborating nodes if needed to assure efficient collaboration. The role of machine intelligence here is not so much to solve the scientific problem but rather to acquire specific data and computational resources that answer recognizable steps in the solution specified by the scientist. The IA complements the scientist who cannot be expected to track the changing holdings and capabilities of a vast computational enterprise.

The sequence described here considers a case requiring new observations. Consequently, the system requests new observations and follow-up production processing. When the on-demand observation is obtained, it is registered with the archive and production processing starts automatically. Eventually, the results are sent to the application tool. Naturally, the scientist would not wait for the return, but within the IA tools can act independently. The application tool may perform an analysis, determine that more data or information is needed and issue additional requests. Eventually all of the available information is assimilated into a "what-if" model and visualization thereof. This result is simultaneously stored in the archive and presented to the scientist.

4. CONCLUSIONS

This paper has introduced the concept of an Intelligent Archive. An IA includes all items stored to support "end-to-end" research and applications scenarios. The stored items could be anywhere in a highly distributed end-to-end system, including on-board caches with the sensors and those in the client systems with the users. Stored items include:

- Data, information, and knowledge
- Software needed to manage holdings
- Interfaces to algorithms and physical resources to support acquisition of data and their transformation into information and knowledge, storing the protocols to interact with other facilities

The phrase IA generalizes the term archive from a simple repository of data to one that supports and facilitates derivation of information and knowledge. An IA has greater ability to operate more autonomously than conventional archives and provides better service to users (as an intelligent assistant) with less operator intervention and, hence, with lower cost. The paper has listed several characteristics of an IA, grouped under three different categories of autonomy: Holdings' Management Autonomy, User Services Autonomy, and System Management Autonomy.

An IA offers new capabilities that distinguish it from archives of today. Some examples of such capabilities are:

- Storing and managing full representations of data, information, and knowledge
- Building intelligence about transformations on data, information, knowledge, and accompanying services involved in a scientific enterprise
- Performing self-analysis to enrich metadata that adds value to the archive's holdings
- Performing change detection to develop trending information
- Interacting as a cooperative node in a "web" of other systems to perform knowledge building (where knowledge building involves the transformations from data to information to knowledge) instead of just data pipelining
- Being aware of other nodes in the knowledge building system (participating in open systems interfaces and protocols for virtualization, and collaborative interoperability)

Many variations and permutations of operational scenarios are possible given the architecture model of elements, component objects, and associations for an intelligent knowledge building system. The conceptual architecture establishes a framework for how future intelligent archives operate in a context of cooperating systems and infrastructures. Furthermore, the architecture is a necessary guide where opportunities for integrating intelligent systems, data understanding, and machine learning can be identified and mapped to goals for automation that increases effective data utilization.

This study of the IA demonstrates the possibility of a new synthesis of ideas and technology. This new synthesis promises unique benefits for science, but they will require more research and effort. Some key technical questions that should be addressed are:

- Throughput requirements of the IDU algorithms in the "context" of an IA
- Appropriate placement of grid technologies within the overall "knowledge building system"
- Likely physical locations (hardware allocations) for the functions of the IA in the light of projected enabling technologies, and the flows of data, information and knowledge. (For example, could a decision support system be hosted on a hand-held device?)
- Design choices regarding integration of e-commerce with science software, for command control vs. peer-to-peer negotiation paradigms, for data- vs. software- mobile paradigms, for micro-sensor data collection, for automated quality assessment, etc.

REFERENCES:

Avery, P. and I. Foster, 2002, http://www.griphyn.org/projinfo/index.php

Cornillon, P., 2000, Personal Communication, December 2000.

Chao, B. F., T. Farr, J. LaBrecque, R. Bindschadler, B. Douglas, E. Rignot, C. K. Shum, J. Wahr, 2002, Understanding Sea Level Changes, IEEE IGARSS and the 24th Canadian Symposium on Remote Sensing, Toronto, Canada, June 24-28, 2002.

DISCOM, 2002, Distance and Distributed Computing and Communication, http://www.cs.sandia.gov/discom/main.html

ESIPFED, 2002, http://www.esipfed.org/

Frew, J., M. Freeston, L. Hill, G. Janée, M. Larsgaard, Q. Zheng, 1999, "Generic Query Metadata for Geospatial Digital Libraries," Proceedings of the Third IEEE META-DATA Conference, http://www.computer.org/proceedings/meta/1999/papers/55/jfrew.htm

Globus, 2002, http://www.globus.org/

M. Moore and D. Lowe, 2001, Leveraging Open Source Development in Large Scale Science Data Management Systems, Proceedings of SPIE, Earth Observing Systems VI, San Diego, CA, August 1-3, 2001, Vol. 4483, pp. 291-300.

M. Moore and D. Lowe, 2002, Providing Rapid Access to EOS Data via Data Pools, Proceedings of SPIE, Earth Observing Systems VII, Seattle, WA, July 7-10, 2002.

McDonald, K. R., R. Suresh, L. Di, 2001, The Data and Information Access Link (DIAL), in Proceedings of the American Meteorological Society (AMS) 81st Annual Meeting, 17th Conference on Interactive Information and Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology, Albuquerque, NM, Jan 14-19, 2001, pp. 166-166.

NASA, 1999. "1999 EOS Reference Handbook: A Guide to NASA's Earth Science Enterprise and the Earth Observing System," M. King and R. Greenstone, editors, NASA Goddard Space Flight Center, http://eos.nasa.gov/

NASA, 2000a, Exploring Our Home Planet, Earth Science Enterprise, Strategic Plan, November 2000, http://www.earth.nasa.gov.

NASA, 2000b, Understanding Earth System Change: NASA's Earth Science Enterprise Research Strategy for 2000-2010, http://www.earth.nasa.gov/visions/researchstrat/ResearchStrat.doc

NASA, 2002a, Earth Science Technology Office (ESTO), http://esto.gsfc.nasa.gov

NASA, 2002b, Advanced Information System Technologies (AIST) Capability and Needs, 2002 Version, http://esto.gsfc.nasa.gov/programs/aist/

NASA, 2002c, Earth Science Enterprise Applications Strategy for 2002-2012, January 2002.

NASA, 2002d, Computing, Information and Communications Technology Program, Intelligent Systems Project, http://is.arc.nasa.gov/

NEESgrid, 2001, http://www.neesgrid.org/

Olsen, L.M., 2000, “Discovering and Using Global Databases,” in Global Environmental Databases: Present Situation; Future Directions, R. Tateishi and D. Hastings, editors, International Society for Photogrammetry and Remote Sensing (ISPRS).

Pfister, R., 2001, “The Information Management System of NASA's EOS Data and Information System,” in Proceedings of the American Meteorological Society (AMS) 81st Annual Meeting, 17th Conference on Interactive Information and Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology, Albuquerque, NM, January 14-19, 2001, pp. 161-163.

PPDG, 2002, http://www.ppdg.net/

Ramapriyan H. K. and G. R. McConaughy, “Version 0 EOSDIS - An Overview,” Technical Papers, 1991 ACSM-ASPRS Annual Convention, Baltimore, MD, March 1991.

Ramapriyan, H. K., S. Kempler, C. Lynnes, G. McConaughy, K. McDonald, R. Kiang, S. Calvo, R. Harberts, L. Roelofs, D. Sun, 2002a, “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, College Park, MD,
April 15-18, 2002.

Ramapriyan, H. K., S. Kempler, C. Lynnes, G. McConaughy, K. McDonald, R. Harberts, L. Roelofs, 2002b, “Conceptual Study of Intelligent Archives of the Future,"NASA Report (in preparation), NASA Goddard Space Flight Center, August 2002.

Ramapriyan, H. K., 2002, Satellite Imagery in Earth Science Applications, Chapter 3 in Image Databases: Search and Retrieval of Digital Imagery, V. Castelli and L. D. Bergman, editors, John Wiley and Sons, NY, pp. 35-82.

Raymond C. A., P. R. Lundgren, S. N. Madsen, J. B. Rundle, 2002, Understanding and Responding to Earthquake Hazards, IEEE IGARSS and the 24th Canadian Symposium on Remote Sensing, Toronto, Canada, June 24-28, 2002.

Schnase, J.L., J. A. Smith, T. J. Stohlgren, S. Graves, C. Trees, 2002, Biological Invasions: A Challenge in Ecological Forecasting, IEEE IGARSS and the 24th Canadian Symposium on Remote Sensing, Toronto, Canada, June 24-28, 2002.

Smarr, L., 1998, The Emerging National Technology Grid: Using High-End Computational Science to Predict the Broad-Based Future, http://www.jacobsschool.ucsd.edu/~lsmarr/talks/San%20Diego/index.htm

Smith, S.G. and R.T. Northcutt, 2000, “Improving Multiple Site Services: The GCMD Proxy Server System Implementation,” EOGEO 2000. Earth Observation (EO) & Geo-Spatial (GEO) Web and Internet Workshop,
April 17-19, 2000.
http://gcmd.nasa.gov/conferences/EOGEO2000/Smith_Northcutt.html

Thigpen, W., 2002, http://www.ipg.nasa.gov/

ACKNOWLEDGEMENT:

This study was supported by the Intelligent Systems Project within NASA’s Computer, Information and Communications Technology Program. The opinions expressed in this paper are those of the authors and do not necessarily reflect NASA’s position.