Design and Implementation of a Prototype Data System for Earth Radiation Budget, Cloud, Aerosol, and Chemistry Data

Bryan A. Baum
and Bruce R. Barkstrom
Atmospheric Sciences Division,
NASA Langley Research Center,
Hampton, Virginia

Abstract

The Earth Observing System (EOS) will collect data from a large number of satellite-borne instruments, beginning later in this decade. To make data accessible to the scientific community, NASA will build an EOS Data and Information System (EOSDIS). As an initial effort to accelerate the development of EOSDIS and to gain experience with such an information system, NASA and other agencies are working on a prototype system called Version 0 (V0). This effort will provide improved access to pre-EOS earth science data throughout the early EOSDIS period. Based on recommendations from the EOSDIS Science Advisory Panel, EOSDIS will have several distributed active archive centers (DAACs). Each DAAC will specialize in particular datasets. This paper describes work at the NASA Langley Research Center's (LaRC) DAAC.

The Version 0 Langley DAAC began archiving and distributing existing datasets pertaining to the earth's radiation budget, clouds, aerosols, and tropospheric chemistry in late 1992. The primary goals of the LaRC V0 effort are the following:

1. Enhance scientific use of existing data;
2. Develop institutional expertise in maintaining and distributing data;
3. Use institutional capability for processing data from previous missions such as the Earth Radiation Budget Experiment and the Stratospheric Aerosol and Gas Experiment to prepare for processing future EOS satellite data;
4. Encourage cooperative interagency and international involvement with datasets and research;
5. Incorporate technological hardware and software advances quickly.

1. Introduction

Humanity's effect on the earth's environment is a critical, timely concern. The government of the United States initiated the Global Change Research Program (GCRP) to develop a more complete understanding of how such changes in the global environment might affect us. This effort is currently trying to assess the current state of the earth's environment and to predict its future state. This is a huge effort, involving both theoretical modeling of various aspects of the earth–atmosphere system and the reduction of vast amounts of data.

The success of the GCRP largely depends on how well we handle the data collected by instruments on the earth's surface, in airplanes, and on satellites, together with the products generated from these data. In support of the GCRP effort, NASA has proposed the Earth Observing System Data and Information System (EOSDIS).

EOSDIS is being designed to archive and distribute much of the data previously collected from regional field measurements, satellite platforms, and other sources. In addition, EOSDIS designers intend to provide support for future research satellites and field measurement programs. This distributed data system will provide access to data from programs sponsored by different government agencies such as NASA, the National Oceanic and Atmospheric Administration (NOAA), and the Department of Energy (DOE). It will also support international programs such as the Tropical Rainfall Measurement Mission (TRMM) and the Global Energy and Water Cycle Experiment (GEWEX).

Among its functions, EOSDIS will generate, archive, and distribute a variety of data products, make it easier to combine data from different sources and provide simple and quick access to data. By making the data from many sources easier to use, the EOSDIS designers hope to encourage more interdisciplinary research and a more global view of the earth's environment. This work will foster interaction among researchers currently working in more narrowly confined disciplines such as oceanography, atmospheric sciences, and geochemistry.

To capitalize on existing centers of scientific expertise and to prevent a single point of failure, the EOS Data Panel recommended that EOSDIS have several distributed active archive centers, or DAACs, across the United States. In response to this recommendation, eight DAACs have been organized, and each

©1993 American Meteorological Society

TABLE 1. EOSDIS distributed active archive centers.

DAAC	Location	Science disciplines
ASF—Alaska SAR facility	Fairbanks, AK	Scanning Aperture Radar (SAR) imagery of polar regions
ORNL—Oak Ridge National Laboratory	Oak Ridge, TN	Field experiments
EDC—EROS Data Center	Sioux Falls, SD	Land processes
GSFC—Goddard Space Flight Center	Greenbelt, MD	Climate, stratospheric chemistry, geophysics, atmospheric dynamics
JPL—Jet Propulsion Laboratory	Pasadena, CA	Physical oceanography
LaRC—Langley Research Center	Hampton, VA	Clouds, radiation budget, aerosols, tropospheric chemistry
MSFC—Marshall Space Flight Center	Huntsville, AL	Hydrological cycle
NSIDC—National Snow and Ice Data Center	Boulder, CO	Polar processes, ice, and oceans

DAAC has distinct responsibilities, as shown in Table 1. The number of DAACs will vary as requirements change over time. Indeed, NASA encourages other data collection and distribution facilities, such as NOAA and the University of Wisconsin, to participate in EOSDIS as affiliated data centers, or ADCs. In much of our discussion below, we will concentrate on the DAAC at NASA's Langley Research Center (LaRC). The LaRC DAAC will archive and distribute datasets related to the earth's radiation budget, clouds, aerosols, and tropospheric chemistry.

2. Version 0 prototyping activity

Although the massive effort for EOSDIS will require several more years of planning and procurement before the scientific community can access its full range of capability, work is already under way to improve data access and to provide experience for guiding the EOSDIS development. A major part of this work starts with existing datasets. Each dataset is currently maintained in its own unique "native" data format. Each DAAC will provide the "native"-formatted data in its area of responsibility. In addition, the LaRC DAAC will make the level 3 (uniform space–time grid) data more accessible to the user community by converting it into a standard data format. We discuss reformatting in section 3b. The impetus for the reformatting effort stems from the past difficulty of the authors of this paper as well as other LaRC scientists in learning how to work with the variety of formats used by various scientific groups. The amount of time spent in figuring out how to decode a binary data stream can be significant since it depends on familiarity with the system that the data came from and the software provided with the data (if such software exists). These disparate collections of data need to have their characteristics documented, and have software developed to enable quick distribution of the data to the users. We call this initial work the EOSDIS Version 0, or V0. The lessons learned from the V0 effort will be provided to the contractors chosen for implementing the next version of the system, Version 1.

The EOSDIS V0 designers will try to provide a uniform interface between the users and the data contained in the various DAACs. Thus, some of the V0 effort is devoted to testing standards that all the DAACs (and ADCs) will have to observe. Likewise, the V0 effort will exercise various kinds of user interfaces and data access methods. The goal of this effort is to improve user data access and provide tools for increasing the user's efficiency. In other words, scientists should be able to find more data that are useful to them, obtain data more quickly than they do now, and use the data more efficiently. This effort requires research into data formats, metadata (information about the data), browsing techniques, and data visualization, as well as more mundane concerns such as simply finding places to store the reformatted data.

The V0 activity provides a unique opportunity to begin testing our ideas about how a distributed data center system needs to work, and is not unique to LaRC. Table 1 shows the distinct responsibilities of each DAAC. In addition, there is a systemwide effort to coordinate all the V0 activities between the DAACs under guidance from the EOSDIS project. The individual DAAC V0 efforts differ at each site according to the type of data and the user community associated with the data. However, it is hoped that even though each DAAC will have different approaches to solving problems, a consolidated approach may develop during the coming year. When this happens, each DAAC will look the same to a user; there will be a similarity in the methods of accessing and ordering data. In the following discussion, we provide a brief overview of the design and implementation of the LaRC Version 0 DAAC.

3. Version 0 design

The LaRC V0 DAAC effort will archive and distribute existing datasets from the Earth Radiation Budget Experiment (ERBE), the Surface Radiation Budget (SRB), the International Satellite Cloud Climatology Project (ISCCP), the Global Tropospheric Experiments (GTE), the Stratospheric Aerosol Measurement (SAM II) experiment, the Stratospheric Aerosol and Gas Experiment (SAGE), and the Measurement of Air Pollution from Satellites (MAPS) experiment. These datasets are currently distributed through the National Space Science Data Center (NSSDC) and/or the NASA Climate Data System (NCDS) at Goddard Space Flight Center (GSFC). However, NSSDC is currently phasing out its involvement in distributing earth science data. In the near future, these data products will no longer be available from NSSDC, and LaRC will assume responsibility for distribution of these datasets in 1993.

To ensure that the system is answering scientists' needs, each DAAC has an oversight committee called a user working group (UWG) that consists of government and university scientists. The UWG is responsible for making sure that the EOSDIS V0 effort remains focused on the requirements of the scientific community.

Each DAAC includes hardware, software, and the people to make sure that the system runs smoothly. Each V0 DAAC consists of three interconnected computer systems: a data archive and distribution system (DADS), an information management system (IMS), and a product generation system (PGS). A majority of the functions may be performed without human intervention. However, there will always be a need for human involvement. For instance, in making tape copies of selected datasets or in making a CD-ROM for data distribution, people are needed to look at a sample of the products to make sure that the data were correctly copied. A more detailed description of these three components is presented in the next section.

a. Data archive and distribution system

The DADS is the subsystem designed to provide automated storage and distribution of the datasets. This system is used to retrieve and subset data that reside on any of the DADS storage media. The DADS can copy a retrieved dataset onto magnetic tape (9-track, 8-mm, 4-mm) or onto an Ethernet-based network for electronic transfer to the scientist's workstation. In addition, the DAAC can distribute a number of prerecorded tapes and compact discs containing the most widely ordered datasets. Until the network throughput increases significantly, we expect that most scientists will want the data mailed to them along with supporting documentation.

The system hardware presently consists of a Unix workstation, a 1-terabyte (TB) magneto-optical jukebox, and a variety of magnetic tape drives. Software residing on the DADS will manage the media use and the automatic placement and retrieval of files to and from the appropriate levels of storage. In addition, software is under development to aid in subsetting the data.

b. Information management system

The IMS handles most of the system's interaction with users. The subsystem's hardware consists of a Unix workstation and a series of magnetic disks. The input/output rate of the disks will dramatically increase by implementing the Redundant Array of Inexpensive Disks (RAID) technology. With the RAID implementation, the I/O transfer rates between the workstation central processing unit and the disks will approach 20 MB s^{-1}, which is much faster than usual disk implementations. The system software consists of a relational database, user-interface software, and other software necessary for such functions as maintaining the system, providing user services, and tracking orders. We anticipate that users will call in from a variety of remote platforms, ranging from character-based personal computers to Unix workstations. Thus, much attention is being given to providing the simplest, most intuitive, and most easily accessible user interfaces. Menus will be uncluttered, have a clear meaning, and contain no unnecessary information.

With the LaRC IMS, a user will have a number of options available for searching and ordering data, browsing through the metadata, performing a search through the use of key words, and searching for data from other DAACs. We have found the time line utility of great use for determining what data are available. The time line graphically depicts the temporal coverage and availability of each dataset currently available in the LaRC V0 DAAC. With the time line, a user can quickly ascertain the temporal coverage of each dataset and the time periods during which combinations of datasets may overlap.

In addition to the time line, there will also be a large amount of on-line documentation of each dataset. This documentation is being implemented through the use of hypertext. Hypertext gives users information about data much as books and catalogs in a library do.

Each V0 DAAC is charged with developing an interface for use by its intended scientific community. In conjunction with the internal development of the individual DAAC interfaces, each with its set of menus and functionality, the V0 project is tying the various requirements of each DAAC together into a synthesized IMS interface that will be standard across all the DAACs. The reason for this effort is twofold. First, the individual DAACs are in various states of development, and many already have some degree of existing functionality. Some DAACs, such as at LaRC, are developing a system for the first time. It is impossible for each DAAC to suddenly arrive at a common set of interfaces so that a user can tie into each individual system and see a similar set of menus. The project has this coordination role. Second, it is the intention of EOSDIS that users will be able to see a common set of interfaces at each of the DAACs. Toward that end, LaRC and the other DAACs are developing their IMS interfaces with the future in mind. The project is intended to have a common IMS interface for all the DAACs in 1994 when V0 is fully operational. In the meantime, each DAAC will be developing an interim set of interfaces for use by the user community.

One important function of a data system is to provide information on available software to help in scientific analysis and visualization of data.

There are a number of existing data systems that have developed interfaces, such as the NASA Climate Data System at GSFC and the Global Land Information System (GLIS) in Sioux Falls, South Dakota. We will take advantage of their expertise during the process of building an interface appropriate for the LaRC datasets and in developing system level IMS expertise.

We are currently investigating a number of issues regarding the use of the IMS, including strategies for metadata searches, shared scientific processing libraries (SPL), standard data format (SDF) software, DAAC-to-DAAC interoperability (searching for and ordering data from other DAACs), ordering information, and user services. The question of whether it is valid to expect that queries may be passed among the eight DAACs has been raised by a number of forums. However, the interoperability concept has been tested among four of the eight DAACs with a fair degree of success. As part of a system-level task initiated by the project office, this particular aspect of the system has been demonstrated to work well when the networks are working. The LaRC DAAC became an operational part of the system-wide set of DAACs in January 1993.

1) Metadata

The ability to maintain an accurate database of information about the data, called the metadata, is almost as important as producing the data itself. Elements of the metadata are the data source and data producers, temporal and spatial extent, resolution, lists of articles published by users of the data, and other pertinent information. Users may query and retrieve this type of information from the database. The metadata for the holdings of each DAAC will also be kept in the Global Change Master Directory (GCMD), located at GSFC. The metadata at a given DAAC will cover only those datasets that are maintained at that DAAC, while the GCMD will incorporate some of the metadata from all the DAACs. A more experienced user will be able to access the data and/or metadata directly at any particular DAAC, while a novice user might be helped best by the GCMD.

Some of the LaRC datasets are currently maintained at the NASA Climate Data System at GSFC. The datasets chosen for distribution at NASA Langley Research Center and the metadata developed for these datasets at GSFC will be moved to LaRC, so there will be a smooth transition between centers responsible for these data.

Metadata are being rewritten during this transition period through discussions between the DAAC and the science groups responsible for producing the data. For previously developed datasets, no further development of metadata is planned at this time except to make sure that the available metadata are correct. For the new data products being produced, such as SAGE and SRB globally gridded level 3 datasets, metadata is provided by the science teams producing the data.

2) Science Processing Library

One important function of a data system is to provide information on available software to help in scientific analysis and visualization of data. We will incorporate into the Science Processing Library (SPL) an array of public domain tools and/or software routines to help expedite research with the data provided by the LaRC DAAC and other DAACs. If there is useful software available from either commercial vendors or other remote institutions, we will provide information on the capabilities of these software packages and how to obtain the software. The SPL will include information on both public domain and commercial software tools as well as software developed and provided by the scientific community.

The SPL will include software for input/output (I/O) functions, browsing, subsetting, imaging, and many other functions. The SPL might also eventually include math and statistics tools, geographic standards tools to aid in incorporating terrain elevation maps and land-use maps, and radiative transfer codes. The success of the SPL will depend on the willingness of the scientific community to share software developed specifically for the analysis of a dataset obtained at a DAAC. The SPL should help data users greatly, and ensure that more time is spent analyzing the data rather than writing software.

We anticipate that the user will want to work with the data at a remote system rather than with DAAC resources. While we expect that there will be a need for individuals to use DAAC resources, such as CPU or resident software, we also expect that most of the analysis will be performed more efficiently at the scientist's institution. The software we provide will be easily transportable to almost any common computer system. Any public domain software that is developed by LaRC DAAC personnel will be continually maintained and updated as the need arises. However, if we provide information on public domain software that is developed and maintained at a remote institution, we cannot guarantee that the software will be free of problems or that the user will be able to influence further development of that software.

Most of our effort with the development of the LaRC V0 DAAC will concentrate on providing quick and relatively simple access to data. Therefore, we will not be placing much emphasis on developing analysis tools. We feel that the development of analysis tools is better provided by educational institutions or commercial vendors.

3) Standard data format

EOSDIS personnel participated in an evaluation of the current standard data formats (SDFs). Samples of existing datasets were reformatted using several different formats. The participants discussed the strengths, problems, and deficiencies of each data format. The list of possible data formats was reduced to three after much discussion. The final candidates were HDF (hierarchical data format) developed by the National Center for Supercomputing Applications (NCSA), netCDF (network common data form) developed by the Unidata Program Center, and CDF (common data format) developed at NASA GSFC. The V0 project chose HDF to be the standard data format. If technical or resource considerations preclude the use of HDF by a DAAC, the data producers may use either netCDF or CDF for data distribution. NCSA will soon merge the capabilities of netCDF and HDF and provide further extensions that will make HDF a more comprehensive standard data format.

NCSA developed HDF to support scientific visualization and analysis software. The designers wanted HDF to be able to incorporate a variety of data models, such as raster images, rectilinear gridded data, irregular mesh data, two- and three-dimensional polygonal data, and text. The HDF software is available through anonymous ftp (file transfer protocol) from the NCSA. A further benefit of HDF is the availability of relatively inexpensive software that enables users to browse HDF files quickly.

4) Interoperability

The main goal of DAAC-to-DAAC interoperability is to allow users to simultaneously query any or all of the DAACs to determine data location and availability. For instance, a user may begin with a gridded product, such as ISCCP C1, and want to focus attention on a particular region. After preliminary investigation, the user could perform a secondary analysis using image data such as that from the Advanced Very High Resolution Radiometer (AVHRR). This type of effort would involve two DAACs, since ISCCP and AVHRR data products will not currently be held at the same site. A user calling into one V0 DAAC will be able to query and order data stored at another DAAC. The system will be transparent to the user, so that all data may be ordered from one session with a DAAC or the system-level IMS. To facilitate the process of inter-DAAC data transfer, V0 will implement a special inter-DAAC network.

5) User and data services

As part of facilitating access to global change data, the V0 effort has formed a User Services Working Group (USWG) to encourage a grass-roots user community to focus on the issues at the various data centers. The USWG is a distinct group and separate from the UWG mentioned in the introduction to section 3.

The USWG is comprised of members from each DAAC who help facilitate user requests and help solve problems. The role of the USWG is to address the needs of the users, a complex task since each user has different requirements. For instance, the image-data user who works with Landsat or AVHRR will order data much differently from a user involved with analyzing gridded, global data products such as ERBE, SAGE, or ISCCP. The system must allow users to explore the metadata in different ways and learn more about possible uses for the data. At the same time, the system must not preclude someone from entering a specific order and bypassing a metadata search. While the LaRC V0 DAAC will be almost a completely automated system, there will be human contacts during normal working hours to help facilitate user requests and to help solve problems. Now that the LaRC Version 0 system is made available to the public, we welcome any comments that will act to ease and facilitate any or all aspects of the Version 0 system.

c. Product generation system

In the EOS time frame, the PGS will be used to convert the raw instrument data received from future satellite platforms into science products. This system must handle the software that geolocates and calibrates raw instrument data before it produces higher-level products. A major effort in this process is to track down and understand the anomalies that often appear in raw satellite data. These anomalies may be due, for example, to a scanner that briefly deviates from its scanning pattern and returns to stowed position, or from an anomaly in the power supply. The role of the PGS is to implement algorithms in the form of computer code supplied by the science teams and to perform functions at the direction of the science teams.

In the V0 time frame, the PGS will work with existing datasets from the IMS and DADS. To learn more about the capabilities that the PGS should have, the LaRC DAAC PGS is planning to begin production of the Surface Radiation Budget products in 1993. These products, described in the next section, are currently being produced at the Satellite Data Archive Center (SDAC) located at LaRC. The SDAC production functions will be merged with the LaRC DAAC during the next year. The PGS will be used to schedule SRB production, import data from the DADS, produce the SRB product, and develop browse products and metadata for ingestion in the IMS. In addition, this exercise will provide experience in making the three subsystems (DADS, IMS, and PGS) work as a single cohesive system.

4. LaRC datasets

The LaRC DAAC initially is responsible for the maintenance and distribution of the following datasets.

a. The Earth Radiation Budget Experiment (ERBE)

ERBE provided the first opportunity to globally measure the earth's radiant energy using identical instruments flying on three separate satellites (Barkstrom 1984; Barkstrom and Smith 1986). Two sets of ERBE sensors flew simultaneously on the *NOAA-9* spacecraft and on the *Earth Radiation Budget Satellite (ERBS)*. A third ERBE instrument set flew on *NOAA-10*, providing 1 month of coverage with all three instrument sets. Since then, these instrument sets have made simultaneous measurements from the *ERBS* and *NOAA-10* instruments. The radiation budget consists of both incident and reflected solar radiation and the longwave radiation emitted to space. The ERBE datasets have been used to explore such topics as the longwave diurnal variability (Harrison et al. 1988) and the role of cloud forcing (Ramanathan et al. 1989) on the earth's radiation budget. The nonscanners on all three satellite platforms continue to operate.

b. Stratospheric Aerosol Measurements (SAM II) and the Stratospheric Aerosol and Gas Experiment (SAGE I & II)

The SAM II sensor (Yue et al. 1984) is a satellite-mounted sun photometer with a single spectral channel centered at 1 μm, and uses a solar occultation or limb extinction technique to obtain aerosol data in polar regions. The solar occultation technique is employed during each sunrise and sunset event encountered by the spacecraft as it orbits the earth. SAM II was launched in October 1978 aboard *Nimbus-7* and is still providing aerosol data.

SAGE I and II (e.g., McCormick et al. 1989a) measure Mie, Rayleigh, and gaseous extinction profiles also using the solar occultation technique. The SAGE I instrument, launched in 1979, operated until 1981 and contained four spectral channels centered at 0.38, 0.45, 0.6, and 1.0 μm. SAGE II added three spectral channels to the four from SAGE I for a total of seven. One channel was added at 0.94 μm for H_2O retrieval and another at 0.525 μm for a better aerosol characterization. The SAGE I 0.45-μm channel was divided into two narrowband channels for the SAGE II instrument in order to improve the retrieval of NO_2.

The SAM II (1978–present), SAGE I (1978–1981), and SAGE II (1984–present) datasets have provided an outstanding long-term quantitative record of volcanic perturbations to the stratosphere. In addition, these data have provided the first reliable measurements of polar stratospheric clouds (e.g., McCormick et al. 1989b) and have provided measurements of stratospheric ozone and upper-tropospheric water vapor.

c. Measurement of Air Pollution from Satellites (MAPS)

The MAPS instrument, which has been flown on two shuttle flights, uses a gas-filter correlation technique to maintain a highly effective spectral resolution in the 4.67-μm fundamental band of CO. This experiment was designed to measure remotely the global distribution of middle- and upper-tropospheric CO (e.g., Reichle et al. 1986; Reichle et al. 1990). This experiment is scheduled to be flown on the Space Radar Laboratory in 1993, 1994, and 1996 aboard shuttle flights.

d. The International Satellite Cloud Climatology Project (ISCCP)

In July 1983, ISCCP began collecting a uniform, global climatology of satellite-measured radiances. ISCCP was chartered to determine the global distribution, variation, and properties of clouds (Schiffer and Rossow 1983). The operational cloud algorithm uses a bispectral technique using a visible channel and a longwave window channel. A large component of ISCCP has been the validation effort to compare ISCCP cloud retrievals with field data, most notably from the First ISCCP Regional Experiment (FIRE) cirrus and stratocumulus field missions held in 1986 (e.g., Starr and Wylie 1990), 1987 (e.g., Minnis et al. 1992), and 1991. The LaRC DAAC will also carry the FIRE datasets.

e. The Global Tropospheric Experiments (GTE)

NASA initiated the GTE to probe atmospheric chemical changes over the Atlantic Ocean, tropical rain forests, and Arctic tundra. GTE collects data through surface- and aircraft-based measurements. However, satellite data have also played an important role, an example being the use of Landsat data to study surface images. In order to investigate current measurement capabilities and foster development of more accurate experimental techniques, scientists and experimenters participated in a series of rigorous intercomparisons called the Chemical Instrumentation Test and Evaluation (CITE) projects. The GTE field measurements were performed using the advanced instruments validated by CITE data. The initial GTE field expeditions, called the Atmospheric Boundary Layer Experiment (ABLE) projects, were designed to investigate the relationships between the biosphere and the atmosphere. The LaRC V0 DAAC anticipates archiving and distributing the data from ABLE-2A and ABLE-2B (Harriss et al. 1988; Harriss et al. 1990) and other missions.

f. Surface Radiation Budget (SRB)

The SRB project, centered at LaRC, was designed to investigate the incoming and outgoing radiation at the earth's surface. The general approach is to use recalibrated ISCCP 3-hourly gridded data supplemented by the ERBE data as input to SRB algorithms that in turn are used to estimate various surface parameters. At this time, two shortwave (SW) algorithms meet accuracy requirements established by the World Climate Research Program (WCRP) (Suttles and Ohring 1986). Each of the two algorithms apply both spectral and angular corrections to the ISCCP data. The first algorithm, developed by Pinker and Laszlo (1992), is iterative in nature and is based on delta–Eddington radiative transfer calculations. The second algorithm, developed by W. F. Staylor (Darnell et al. 1988), uses a parameterized physical model that incorporates optical physics relations to calculate cloud and aerosol transmission. The equations have been tuned with historical data from around the globe. Initial SW SRB products include, for example, all-sky net surface-absorbed flux, surface "cloud forcing" flux, top-of-atmosphere net flux, and atmosphere absorbed flux. Ground truth for downward surface flux, courtesy of the Swiss Federal Institute of Technology, is also furnished for verification of the satellite-derived products.

5. Concluding remarks

The full implementation of EOSDIS involves a host of issues not previously explored in scientific data systems. Specifically, the success of EOSDIS requires that a series of distributed data centers work closely together in order to provide full access to the past, current, and future science data products. Each DAAC must "look" the same to the user; each DAAC must provide a range of services required by the users. These services will depend to some degree upon the type of data being processed and archived at each center.

The Version 0 prototyping effort provides the first evidence of how the EOSDIS will ultimately work. Each of the designated DAACs is currently combining existing data systems into a central data system, and at the same time working closely with the other DAACs in order to ensure that the interfaces between DAACs will be seamless. That is, a user will have to learn only a single protocol in order to access any of the data centers.

The Langley Research Center, by the nature of its expertise in the fields of the earth's radiation budget (ERBE), aerosols (SAGE), and chemistry (GTE, MAPS), has been designated as the processing center for future experiments in these fields. As part of the Version 0 work, LaRC is now preparing to support these datasets in an on-line system. The system underwent thorough testing from July through December 1992, and is now ready for public use.

A cost structure for the data has yet to be implemented; costs are to be minimal. The cost of the data to be ordered will appear on a screen before the order is placed, and the user will have ample opportunity to adjust the order according to the cost. The computer time used to access and order the data will be free; however, tape costs must be reimbursed by sending replacement data-grade magnetic tapes of equal or greater quality.

We believe that the LaRC V0 DAAC will provide data to the researcher with a minimum of difficulty. We hope that use of our system will encourage more interdisciplinary research. Researchers interested in accessing this system should contact our user services staff by sending e-mail to `userservices@eosdis.larc.nasa.gov`. Further information regarding these matters may be obtained by writing to the above e-mail address or contacting the authors.

References

Barkstrom, B. R., 1984: The Earth Radiation Budget Experiment (ERBE). *Bull. Amer. Meteor. Soc.,* **65,** 1170–1185.

——, and G. L. Smith, 1986: The Earth Radiation Budget Experiment: Science and implementation. *Rev. Geophys.,* **24,** 379–390.

Darnell, W. L., W. F. Staylor, S. K. Gupta, and F. M. Denn, 1988: Estimation of surface insolation using sun-synchronous satellite data. *J. Climate,* **1,** 820–835.

Harrison, E. F., D. R. Brooks, P. Minnis, B. A. Wielicki, W. F. Staylor, G. G. Gibson, D. F. Young, F. M. Denn, and the ERBE Science Team, 1988: First estimates of the diurnal variation of longwave radiation from the multiple-satellite Earth Radiation Budget Experiment (ERBE). *Bull. Amer. Meteor. Soc.,* **69,** 1144–1151.

Harriss, R. C., S. C. Wofsy, M. Garstang, E. V. Browell, L. C. B. Molion, R. J. McNeal, J. M. Hoell, Jr., R. J. Bendura, S. M. Beck, R. L. Navarro, J. T. Riley, and R. L. Snell, 1988: The Amazon Boundary Layer Experiment (ABLE-2A): Dry season 1985. *J. Geophys. Res.,* **93,** 1351–1360.

——, M. Garstang, S. C. Wofsy, S. M. Beck, R. J. Bendura, J. R. B. Coelho, J. W. Drewry, J. M. Hoell, Jr., P. A. Matson, R. J. McNeal, L. C. B. Molion, R. L. Navarro, V. Rabine, and R. L. Snell, 1990: The Amazon Boundary Layer Experiment: Wet season 1987. *J. Geophys. Res.,* **95,** 16 721–16 736.

McCormick, M. P., J. M. Zawodny, R. E. Veiga, and J. C. Larsen, 1989a: An overview of SAGE I and SAGE II ozone measurements. *Planet. Space Sci.,* **37,** 1567–1586.

McCormick, M. P., C. R. Trepte, and M. C. Pitts, 1989b: Persistence of polar stratospheric clouds in the southern polar region. *J. Geophys. Res.,* **94,** 11 241–11 251.

Minnis, P., P. W. Heck, D. F. Young, C. W. Fairall, and J. B. Snider, 1992: Stratocumulus cloud properties derived from simultaneous satellite and island-based instrumentation during FIRE. *J. Appl. Meteor.,* **31,** 317–339.

Pinker, R. T., and I. Laszlo, 1992: Modeling surface solar irradiance for satellite applications on a global scale. *J. Appl. Meteor.,* **31,** 194–211.

Ramanathan, V., B. R. Barkstrom, and E. F. Harrison, 1989: Climate and the earth's radiation budget. *Phys. Today,* **42,** 22–32.

Reichle, H. G., Jr., V. S. Connors, J. A. Holland, W. D. Hypes, H. A. Wallio, J. C. Casas, B. B. Gormsen, M. S. Saylor, and W. D. Hesketh, 1986: Middle and upper tropospheric carbon monoxide mixing ratios as measured by a satellite-borne remote sensor during November 1981. *J. Geophys. Res.,* **91,** 10 865–10 887.

——, ——, ——, R. T. Sherrill, H. A. Wallio, J. C. Casas, E. P. Condon, B. B. Gormsen, and W. Seiler, 1990: The distribution of middle tropospheric carbon monoxide during early October 1984. *J. Geophys. Res.,* **95,** 9845–9856.

Schiffer, R. A., and W. B. Rossow, 1983: The International Satellite Cloud Climatology Project (ISCCP): The first project of the World Climate Research Programme. *Bull. Amer. Meteor. Soc.,* **64,** 779–784.

Starr, D. O'C., and D. P. Wylie, 1990: The 27–28 October 1986 FIRE IFO cirrus case study: Meteorology and clouds. *Mon. Wea. Rev.,* **118,** 2259–2287.

Suttles, J. T., and G. Ohring, 1986: Surface Radiation Budget for climate applications. NASA Reference Publication 1169, 132 pp.

Yue, G. K., M. P. McCormick, and W. P. Chu, 1984: A comparative study of aerosol extinction measurements made by the SAM II and SAGE satellite experiments. *J. Geophys. Res.,* **89,** 5321–5327.