


NVODS AND THE DEVELOPMENT OF OPeNDAP

BY PETER CORNILLON, JENNIFER ADAMS,
M. BENNO BLUMENTHAL, ERIC CHASSIGNET, ETHAN DAVIS,
STEVE HANKIN, JAMES KINTER, ROY MENDELSSOHN,
JAMES T. POTEIRA, ASHWANTH SRINIVASAN, AND JOSEPH SIROTT



ABSTRACT. The National Oceanographic Partnership Program (NOPP) funded a project to develop the foundation for a National Virtual Ocean Data System (NVO DS) that has resulted in a robust data access framework for the exchange of oceanographic data (the Open source Project for a Network Data Access Protocol, or OPeNDAP) and a broad community of ocean data providers that remains vigorous and growing five years after NOPP funding ended. The project produced a number of “lessons learned” related to the design and implementation of distributed data systems that can inform other related efforts. These lessons are presented along with a brief overview of OPeNDAP and summaries of a number of projects that depend on OPeNDAP for data distribution.

INTRODUCTION

Following the 1999 National Oceanographic Partnership Program (NOPP) Broad Agency Announcement, a consortium of 51 private, academic, state, federal, and international partners and collaborators were funded to plan and implement a network-based system—the National Virtual Ocean Data System (NVO DS)—to provide for the discovery of and seamless access to oceanographic data. The core infrastructure of the system—the component providing the “glue” among system elements—was based on OPeNDAP (the Open source Project for a Network Data Access Protocol), referred to as the Distributed Oceanographic Data System (DODS) at the time. The project team was divided into four regional groups (the Northeast, the Southeast, the Gulf Coast, and the West Coast¹) with a coordinator for each group (Linda Mercer of the Department of Marine Resources, State of Maine; Anne Ball of the Coastal Services Center, National Oceanic and Atmospheric Administration [NOAA]; Worth Nowlin of the College of Geosciences, Texas A&M University; and Mark Abbott of the College of Oceanic and Atmospheric Sciences, Oregon State University; respectively). Each coordinator organized one or more regional meetings to

obtain input from the community about what data should be the focus of the effort in their community and how the data access protocol should be modified to best meet overall project objectives. Following the regional meetings, two national meetings were held to synthesize the recommendations emerging from the regions. Based on the recommendations of the more than 100 oceanographers, coastal planners/managers, data curators, and data system administrators participating in the regional and national workshops, the data access protocol was refined, data servers were installed, and pilot projects were initiated.

NVO DS as an integrated data system ended with the end of NOPP funding. However, OPeNDAP, the core infrastructure developed for NVO DS, has proven to be a robust data access framework for the exchange of data that continues to see significant use in the oceanographic community. The focus of this manuscript is on this infrastructure and how it is being used today. Specifically, we summarize the use of OPeNDAP at nine institutions, in some cases to serve their data and in others to Web-enable data access tools that they have developed. Although many other organizations use the protocol either to share data internally or to provide public access to

their data holdings, the nine institutions highlighted herein are representative of the OPeNDAP data provider community in general. Although the focus in this manuscript is on data providers, it is important to keep in mind that there are thousands of users of data provided via OPeNDAP.

In the remainder of this section, we provide a brief overview of OPeNDAP. Following this section, nine sample implementations are described, and the final section presents lessons learned from the combined efforts of these and other implementations of the protocol.

The History of OPeNDAP

In the early 1990s, it became clear that there was a need for a data access protocol that would facilitate network access to principal-investigator-held data as well as to data held in federal archives. For example, approximately 90% of the satellite-derived data (≈ 10 MB/day) acquired by the remote-sensing group at the University of Rhode Island was being obtained from scientific colleagues, not from federal archives, and yet the data systems being envisioned at the time were being designed for access to data from large federal archives. In light of this mismatch between the design of data systems and the use patterns within the scientific community, a workshop was held at the University of Rhode Island to design and implement a data access protocol that would work equally well for the principal investigator as for federal archives (Cornillon et al., 1993). The system defined at this workshop was referred to as DODS, and the first implementation of its data access protocol was released in 1995. The National Aeronautics and Space Administration

¹ A Great Lakes Region was envisioned in the original proposal, but dropped in the middle of the first year because of the lack of a coordinator.

(NASA) and NOAA funded this effort.

The fundamental constraints imposed on the design of the DODS data access protocol and associated data servers and clients were to:

- minimize the impact on data providers of making their data accessible via the network
- isolate the format in which the data were stored from the data user
- allow data users to access the data directly from their application packages
- allow data users to acquire only those data of interest (i.e., allow the user to request data subsets from the server)

Also important in the design of the system was that the data access protocol was conceived as a discipline-neutral layer that would lie above the network transport protocol. However, the design team recognized that data semantics are critical for their use. The protocol, discussed briefly below, was therefore designed to permit inclusion of semantic information (also referred to as meta-data) with very few constraints on the

structure or content of these semantics, thus allowing those in different disciplines to layer the semantics relevant for their discipline within the data transfers. This is precisely how many in the netCDF (<http://www.unidata.ucar.edu/software/netcdf/>) community, and other communities such as the National Center for Atmospheric Research High-Altitude Observatory, use this protocol. The discipline neutrality of the protocol was important in that this allowed it to be used over a broad range of disciplines, with contributions from other disciplines feeding back to the overall implementation of the core servers and clients. Although not presented as a sample implementation in the next section—those examples are reserved for oceanographic and meteorological applications—the synergy between the use made of the system by the High-Altitude Observatory and the core DODS (now OPeNDAP) group is a good example of this feedback. Specifically, the most recent OPeNDAP server, Hyrax, is based on work undertaken by the High-

Altitude Observatory as part of the Earth System Grid project. The notion of discipline neutrality is of particular importance to the oceanographic community because it encompasses a broad range of subdisciplines (physics, biology, chemistry, and geology), with often quite different semantic (discipline-specific) requirements, and a range of applications (research, commercial, and management) whose semantic requirements also vary.

In 2000, shortly after receiving NOPP funding, a nonprofit corporation was formed to maintain, promote, and evolve the DODS data access protocol. This corporation was named OPeNDAP².

NVODS built on the early DODS work, enhancing the data access protocol (DAP), collaborating with Global Change Master Directory staff to provide an OPeNDAP portal in the Global Change Master Directory for data discovery, installing OPeNDAP servers at a number of oceanographic institutions and federal laboratories, and developing additional OPeNDAP clients.

Peter Cornillon (pcornillon@gso.uri.edu) is Professor of Physical Oceanography, Graduate School of Oceanography, University of Rhode Island, Narragansett, RI, USA. **Jennifer Adams** is Research Scientist, Center for Ocean-Land-Atmosphere Studies (COLA), Institute of Global Environment and Society (IGES), Calverton, MD, USA. **M. Benno Blumenthal** is Data Library Manager, International Research Institute for Climate and Society, Columbia University, Palisades, NY, USA. **Eric Chassignet** is Professor and Director, Center for Ocean-Atmospheric Prediction Studies, Florida State University, Tallahassee, FL, USA. **Ethan Davis** is Software Engineer, Unidata, University Corporation for Atmospheric Research, Boulder, CO, USA. **Steve Hankin** is Research Scientist, Pacific Marine Environmental Laboratory (PMEL), National Oceanic and Atmospheric Administration (NOAA), Seattle WA, USA. **James Kinter** is Director and Senior Research Scientist, COLA, IGES, Calverton, MD, USA. **Roy Mendelssohn** is Supervisory Research Operations Analyst, NOAA/National Marine Fisheries Service/Southwest Fisheries Science Center/Environmental Research Division, Pacific Grove, CA, USA. **James T. Potemra** is Assistant Manager, Asia-Pacific Data-Research Center and Faculty, School of Ocean and Earth Science and Technology, University of Hawaii, Honolulu, HI, USA. **Ashwanth Srinivasan** is Research Assistant Professor, Rosenstiel School of Marine and Atmospheric Science, University of Miami, Miami, FL, USA. **Joseph Sirott** is a software engineer at PMEL, NOAA, Seattle WA, USA.

² The software based on the OPeNDAP data access protocol is often referred to as OPeNDAP software, which can be confusing. For example, an OPeNDAP-compliant server, such as the PyDAP server, may be referred to as an OPeNDAP server even though it was not developed by the OPeNDAP corporation. Therefore, in this article, we use the somewhat more cumbersome "OPeNDAP corporation" when referring to the corporate entity, an "OPeNDAP server/client/..." when referring to software based on the OPeNDAP data access protocol, simply "OPeNDAP" when the type of software is obvious, and "DAP" when referring explicitly to the data access protocol.

The project also maintained a list of data sets being served via DAP. This list is accessible at: <http://www.opendap.org/data/>, but it is now substantially out of date; it simply became too labor intensive to manually maintain the list as the number of data providers and data sets increased. There are currently in excess of 100 institutions hosting OPeNDAP servers, providing network access to several petabytes of data held in thousands of data sets.

As indicated above, it is DAP that provided the “glue” for NVOADS, connecting the user’s analysis software with the data provider’s server. In addition to providing this connection, DAP provides a means for the client application to specify simple projection and subsetting operations in a data request to the server. These operations are applied to the data at the server, and the requested data are returned to the requesting client application. Thus, use of DAP effectively isolates the format in which the data are stored from their use in the client application while allowing for subsetting and projection operations on the data. Example: a user requests a subset of data from a remote archive via a URL. The server extracts the subset of interest, transforms this subset to the OPeNDAP data model, compresses the resulting data object, and sends it to the requesting client. On receipt, the client decompresses the data stream, translates it from the OPeNDAP data model to the data model used by the client application, and enters it into the application’s workspace. The subsetting and projection operations are completely generic, hence, reference to DAP as “discipline neutral.” It is the protocol—DAP—that permits the connectivity required for this example.

PROJECTS/ORGANIZATIONS

In this section, we summarize a number of active projects for which DAP is a key element, and that have resulted or benefited from NOPP NVOADS funding. Some of these projects involve continuations of partnerships begun with the original NOPP award. Others involve groups partnering with Integrated Ocean Observing System (IOOS) funding; OPeNDAP involvement in the IOOS data system is a direct outgrowth of the

original NOPP funding. Finally, several of the projects, although not formally related to either NOPP NVOADS funding or IOOS, are presented because they demonstrate the breadth of OPeNDAP penetration into the meteorological and/or the oceanographic communities—penetration enabled by the NVOADS project. These relationships are shown diagrammatically in Figure 1. Of interest in this figure is the level of project inter-relatedness. These connections are true




Figure 1. Partnerships among the groups discussed in this manuscript. Light green background indicates focus on a client application (see section on client-side applications), light red on a server application (see section on servers), light grey on a data provider (see section on large data archives), and black on the OPeNDAP corporation (see section on OPeNDAP history). Thick blue lines indicate partnerships that began with NOPP National Virtual Ocean Data System (NVOADS) funding and/or “follow-on” NASA Research, Education, and Applications Solutions Network funding, thick red lines indicate partnerships funded under the Integrated Ocean Observing System (also viewed as a follow-on to NVOADS), and thin black lines indicate partnerships that rely on the Data Access Protocol (DAP) but were funded from other sources.

not only of the projects discussed here but also of many other projects making use of DAP; they are a natural outgrowth of building a data system based on open-source software in a highly distributed environment. Specifically, because DAP is an open-source protocol, growth in its use depends on community involvement. There are two fundamentally different contributions that those in the community can make: (1) software developers can develop applications that facilitate access to data via DAP and (2) data providers can make their data accessible via the protocol. Furthermore, software contributions facilitating access to data can be made on the client side or on the server side. Projects/partnerships discussed in the following are divided into these three categories: client-side software contributions, server-side software contributions, and data contributions. Each example is meant to highlight a specific function—indicated by the tag line in the paragraph heading.

In addition to the interrelatedness of the projects, it is interesting to note that although several of the projects involve implementation of complete systems—client, server, and data that work efficiently together—because they rely on DAP for the exchange of data between client and server, the data may be accessed by clients from other systems. Box 1 briefly presents this form of access in the context of the International Research Institute for Climate and Society (IRI) Data Library.

Client-side Applications Facilitating Access to Data Via OPeNDAP

At the lowest level, OPeNDAP data requests are defined by a URL that specifies the data server, the data object of interest at that server, and the subset of this data object to be returned. For arrays, this subset is specified indicially (i.e., by the ranges of each index in the array that together define the region of interest). At best, these URLs can be difficult to construct. However, for data with rich and consistent metadata, once located, constructing the URL can be straightforward. Many of the netCDF-based OPeNDAP-enabled clients, for example, Ferret³, GrADS⁴ (Grid Analysis and Display System), IDV⁵ (Integrated Data Viewer), and EDC⁶ (Environmental Data Connector), take advantage of the fact that a large fraction of netCDF data sets available in meteorology and oceanography use as metadata conventions either the Cooperative Ocean-Atmosphere Research Data Standard (COARDS) or NetCDF Climate and Forecast Metadata Conventions (CF) to facilitate access. However, these clients do not address the often complicated organization of data objects in large gridded data sets frequently associated with numerical model output or satellite-derived archives, nor do they address access for users via Web browsers. In this section, we discuss the two clients of partnerships formed as part of the initial NVOODS funding that address

these issues. In both cases, the user is, to a large extent, isolated from the data structure, not just the format in which the data are stored (the job of DAP). The metadata are augmented where missing or incomplete to facilitate data use, and simple conversions are applied to the data such as converting missing value flags, which tend to differ from one data set to another, to a common value. Unlike the clients cited above, or others such as CDAT⁷, ODC⁸, and EDC⁹, which are based on other metadata conventions, the clients described here will only operate with data sets for which they have been configured.

The Live Access Server: Providing Browser Visualizations for Distributed Data Sets

Within the NVOODS collaboration, the Live Access Server (LAS) project at NOAA's Pacific Marine Environmental Laboratory (PMEL) was the designated provider of a "generic online browse capability"—the capability to provide scientific visualizations via a Web browser. With its initial release in 1993, LAS was one of the first servers on the World Wide Web with the ability to deliver on-the-fly, custom, scientific visualizations (Sirott et al., 2001). Through the NVOODS effort, these visualization services were adapted for remotely hosted OPeNDAP-served data sets.

Using the NVOODS LAS¹⁰, users can request common analysis operations—averages, variances, and

³ <http://ferret.pmel.noaa.gov/Ferret/documentation>

⁴ <http://grads.iges.org/grads/index.html>

⁵ <http://www.unidata.ucar.edu/software/idv/>

⁶ <http://www.asascience.com/software/arcgistools/edc.shtml>

⁷ <http://www2-pcmdi.llnl.gov/cdat>

⁸ <http://opendap.org/ODC/>

⁹ <http://www.asascience.com/software/arcgistools/>

¹⁰ <http://ferret.pmel.noaa.gov/nvoods> – provided by the LAS development group at PMEL as a community service

BOX 1: The International Research Institute for Climate and Society Data Library: An Integrated, Distributed Data Library Based on the Data Access Protocol

Ingrid is the software suite used to implement the International Research Institute for Climate and Society (IRI) Data Library (<http://ingrid.ldeo.columbia.edu/>). It has three major goals: helping users find suitable data sets, providing a useful summary of the data set once found, and helping users make use of the data in that data set. Ingrid uses an object-oriented approach to create a Web-accessible environment where data sets can be found, accessed, viewed, and manipulated.

The IRI Data Library is organized as a library, a collection of both near and far-off data sets, designed to make the data more accessible for the library's users. Data sets in the library come from many different sources and many different "data cultures" in many different formats. By "data set" we mean a collection of data organized as multidimensional dependent variables, independent variables, and sub-data sets, along with the metadata (particularly use metadata) that makes it possible to interpret the data in a meaningful manner.

Ingrid, which provides the infrastructure for the Data Library, is an environment that allows working with data sets, for example, read, write, request, serve, view, select, calculate, and transform (Figure B1). It hides an extraordinary amount of technical detail from the user, letting the user think in terms of manipulations of data sets rather than manipulations of files of numbers. Among other things, this hidden technical detail could be accessing data on servers in other places, doing only the small needed portion of an enormous calculation, or translating to and from a variety of formats and between "data cultures."

Although the data set environment appears to be uniform, simplifying tasks for the user, the data sets are spread over a network of data servers, and local copies are updated automatically. This updating is achieved by using OPeNDAP to exchange data between different Ingrid servers, to get data from other OPeNDAP servers, and to send data to OPeNDAP clients. Although it uses an alternate implementation of the OPeNDAP core that allows selective partial execution and data streaming, there have been no extensions to the OPeNDAP protocol: all the uniquely Ingrid information is encoded in the filename part of the URL, which is, by definition, site dependent.

The Data Library holds primarily oceanic and atmospheric data as well as model runs, with daily to monthly resolution extending decades to centuries. Both gridded and station data are included. It also includes topographic and feature data, for example, administrative boundaries ranging from countries to districts, some health/economic data, and paleoclimate data.

Because the IRI Data Library exchanges data via DAP, all of its data holdings are accessible from DAP-enabled clients. For example, GrADS, a DAP-enabled client discussed in the section on servers, operates most efficiently when used in conjunction with GDS, but it can also access data from the IRI Data Library.

IRI Climate Data Library Overview

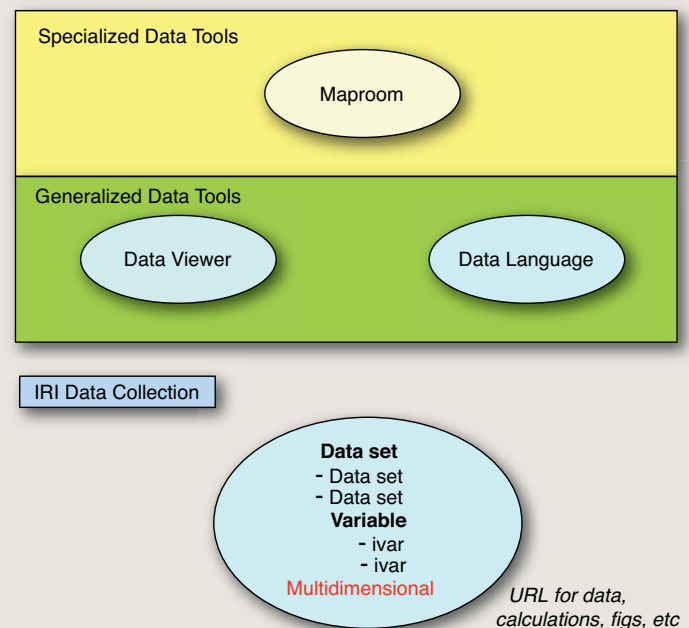


Figure B1. The IRI Data Library is based on making a diverse collection of data appear as a set of nested data sets ultimately containing multidimensional variables. This base layer can deliver data in multiple formats, as images, or navigational pages, and it allows analyses to be performed, creating new data sets on the fly. Generalized data tools built on that fundamental engine include a general data viewer and a data processing language. On top of that are specialized tools that we gather into maprooms, which pick particular data sets and views/manipulations for particular audiences who would not necessarily be able to decide which data set is appropriate for a particular use.




Figure 2. Difference between two time averages of potential temperature at 300 m from an Intergovernmental Panel on Climate Change model run at the Geophysical Fluid Dynamics Laboratory. Upper left: average for 1861–1871. Lower left: average for 1990–2000. Right: average for 1990–2000 minus 1861–1871.

extrema—computed over arbitrary geographical areas, volumes, or time ranges—and visualize the results. Users can compute and visualize differences between remotely hosted OPeNDAP data sets, with LAS regrid-
ding the data as necessary in order to perform the subtraction. Figure 2 shows the image created when averages of potential temperature obtained from a Geophysical Fluid Dynamics Laboratory Intergovernmental Panel on Climate Change (GFDL IPCC) model run for two different periods are differenced.

The NVODS LAS serves products

created on the fly from approximately 120 OPeNDAP-accessible data sets hosted at ten independent institutions distributed across the community. These data sets include ocean, atmosphere, operational, and climate model runs; satellite observations; climatological products; and bathymetries.

The MATLAB OPeNDAP Ocean Toolbox: Exposing the Structure of Complex Data Sets

Shortly before the end of the NVODS project, the University of Rhode Island, with OPeNDAP, received NASA funding

as part of the Research, Education, and Applications Solutions Network (REASoN) program to continue work begun during NVODS but with a narrower focus—satellite-derived surface ocean properties. As part of this funding, together with funding from the NOPP-funded HYbrid Coordinate Ocean Model (HYCOM) project (see later section on HYCOM), the University of Rhode Island group developed a suite of MATLAB Graphical User Interfaces (GUIs)¹¹, collectively referred to as the MATLAB OPeNDAP Ocean Toolbox, to access HYCOM model output and data from

¹¹ <http://oceanographicdata.org/toolbox>

major archives of satellite-derived sea surface temperature, sea surface height, ocean color, and ocean surface winds.

The GUIs are all structured similarly (Figure 3). Each has a section for data set selection if the GUI accesses more than one data set, for variable selection, and for space and time selections, including the ability to select based on resolution, and a section in which the user can select the desired form of data output. All GUIs provide an option for specifying a common structure for the returned data and metadata. This option means that users can develop MATLAB scripts based on this structure and the scripts can be used on the output from any of the GUIs. The GUIs provide another important function—they “expose” the data set in a simple graphical interface. A number of the data sets accessed consist of many files in a complex directory structure. Finding the subset of data of interest in these archives can be very difficult. The suite of GUIs also comes with a front end for finding data of interest within the suite. Finally, the entire suite may be compiled, which means that the GUI may be run on computers without a MATLAB license. The accessed data are output as CF-compliant netCDF files on the user’s computer in this case—files that are readily ingested by many application packages.

Significant effort was devoted to presentation of the archives in the GUI—synthesizing the documentation for the site, understanding the data’s organizational structure, and then presenting this information graphically for the user. Taken together, these functions greatly facilitate use of the data.




Figure 3. The HYbrid Coordinate Ocean Model (HYCOM) MATLAB Graphical User Interface annotated by selection region.

Servers Facilitating Access to Data via OPeNDAP

DAP prescribes the basic operation of OPeNDAP servers—the form of the commands that they accept and the structure of the returned data. A number of low-level servers have been built for a variety of commonly used data formats and in several different programming languages. In this section, we discuss three servers that provide functionality to address specific needs beyond those provided by the basic server.

Dapper: Facilitating Access to In Situ Data Sets

It can be challenging to manage in situ oceanographic data. These data are typically collected from a diverse set of

research groups using different instruments for measurements (including fixed moorings, expendable bathythermographs, conductivity-temperature-depth probes, and profiling floats). They are often archived in a variety of formats. Conventional data models that were developed for data on regular grids (such as satellite or model data) do not work well with in situ data. Furthermore, it is difficult to represent these data using conventional file-based formats because these formats were designed primarily as archive formats for gridded data. Dapper¹² was designed as part of the NVOADS program to address these issues. Specifically, it was designed to provide efficient remote access to large, in situ oceanographic data collections (millions

¹² www.epic.noaa.gov/epic/software/dapper/

of profiles) to allow users to query the collection for data in a particular spatial or temporal region, and to return data to users or applications using a single format (Sirott et al., 2004).

Dapper is based on a two-stage protocol for requesting in situ data. In the first stage, a Dapper-enabled application queries Dapper for the availability of data in a specific geospatial region. Dapper returns (via DAP) a set of unique IDs that correspond to each profile or time series within the region. The application can then use all or part of these IDs to access the physical measurements (again via DAP).

Data are now distributed to thousands of users from several large, in situ data sets, including the World Ocean Database 2005 (several million profiles) and near-real-time data from the Argo profiling floats program (several hundred thousand profiles and growing). Implementations of the Dapper client are available for popular visualization and analysis tools (such as MATLAB, ncBrowse, and GrADS). Finally, the two-stage protocol has proven sufficiently flexible to use as the primary in situ data access source for PMEL's Web-based visualization server, DChart¹³.

THREDDS Data Server: Data Set Aggregation

Unidata has been using OPeNDAP for remote data access for close to ten years with DAP now integrated into several of the Unidata software packages and widely used by the Unidata community. Of particular interest here is the THREDDS (Thematic

Realtime Environmental Distributed Data Services) Data Server (TDS)¹⁴, an OPeNDAP server designed to provide an aggregated view of large, multi-file archives. As mentioned earlier, the organizational structure of large data sets can be complex, consisting of many files residing in a number of directories. This is particularly true of model output and of satellite-derived data sets. In addition to rendering it difficult to navigate through such archives, requesting data from them can be tedious if the user is interested in a small piece of data from each of a large number of files. Suppose that a user wants to access only the data in an area around Bermuda from a 10,000-image archive of North Atlantic sea surface temperature fields. If each two-dimensional (latitude-longitude) image is stored in a unique data file, accessing those data will require 10,000 requests. To define an aggregate data set, TDS makes use of THREDDS machine-readable catalogs of local file locations and/or an OPeNDAP data set URL. With aggregation, these 10,000 two-dimensional images will look like a single three-dimensional (latitude-longitude-time) cube, greatly simplifying access to subsets of the data and, in this case, requiring a single request to access the desired subset of data. In addition, TDS, using Unidata's netCDF-Java library¹⁵, can serve any of the data formats read by netCDF-Java, including netCDF, HDF5 (Hierarchical Data Format-5), GRIB-1 (Grid in Binary-1), GRIB-2, GINI (GOES Ingest and NOAA/PORT Interface), and GEMPACK (General Equilibrium

Modelling PACKage). TDS is also a multi-protocol server, providing Web Coverage Service (WCS)¹⁶ access to OPeNDAP-enabled data sets.

The GrADS Data Server: Server-Side Processing

The Center for Ocean-Land-Atmosphere Studies (COLA) developed GrADS as an interactive desktop tool for accessing, manipulating, and visualizing earth science data. In the late 1990s, GrADS was modified with DAP so that it could serve as either a client, GrADS, or a server, the GrADS Data Server (GDS). The combination of OPeNDAP technology and GrADS capabilities has dramatically altered the way that COLA organizes and interacts with data; users can access, subset, analyze, and display data that reside at a remote location, without transporting the data used and without any detailed knowledge of the data's structure or format.

The core of GDS is OPeNDAP software, but with several essential enhancements. Services are provided for all GrADS-readable data formats: GRIB, binary, netCDF, HDF, BUFR (Binary Universal Format Representation), and GrADS in situ data format. GDS unifies all these data formats into a single netCDF framework. GDS also provides a server-side analysis capability—based on permissible GrADS operations—that allows the user to perform operations on one or more data sets on the server. Examples of analysis operations are averaging, smoothing, differencing, correlation, and regression.

Any DAP-enabled client can access

¹³ <http://www.epic.noaa.gov/epic/software/dchart/>

¹⁴ <http://www.unidata.ucar.edu/projects/THREDDS/>

¹⁵ <http://www.unidata.ucar.edu/software/netcdf-java/> and <http://www.unidata.ucar.edu/software/netcdf/>

¹⁶ WCS is a data access protocol for gridded data developed by the OGC

GDS, and GrADS, as a client, and can access any DAP-enabled server, as discussed in Box 1. But, the GrADS/GDS client/server combination is especially powerful because the structure and metadata of any GDS data set are completely understood by GrADS; thus, no additional metadata are required. In addition, the server-side analysis expressions accepted by GDS are the same as those handled by GrADS, so GrADS users are well positioned to exploit this capability.

Some Large Data Archives Being Served via DAP

There are a number of large data archives being served via DAP. In this section, we profile three archives to provide the reader with a sense for some of these. We also briefly summarize the holdings of Unidata and COLA, two groups discussed previously in the context of software contributions, in that they are representative of many software development efforts that have been undertaken by data providers in an effort to better serve their user communities.

SWFSC: The Data Passthrough

The Southwest Fisheries Science Center's (SWFSC) Environmental Research Division (ERD), through the combination of THREDDS and OPeNDAP services, provides Internet access to a broad range of high-level physical and biological data encompassing more than 20 TB of data representing approximately 300 different data sets. Of particular interest is that ERD has taken advantage of DAP's ability to "serve"

remote data to act as a portal to other NOAA DAP-enabled data providers with similar holdings (National Climatic Data Center, National Coastal Data Development Center, PMEL, and other National Marine Fisheries Service labs) as well as to non-NOAA data providers, such as the US Geological Survey. Specifically, in addition to aggregating and subsetting the data being served, TDS can be used to "serve" and aggregate data from other (remote) OPeNDAP or THREDDS servers, providing the user with a much larger virtual catalog of data without having to copy and update the original data. A significant subset of the data ERD serves is being served in this fashion.

HYCOM: Access to Very Large Archives of Numerical Model Output

The HYCOM consortium is a partnership among universities, government laboratories, and private industry engaged in developing high-resolution, eddy-resolving, global ocean prediction systems (Chassignet et al., this issue). Ocean prediction systems produce estimates and forecasts of the basic oceanic state, including currents and associated flow features such as eddies, fronts and jets, water temperature, and salinity. Reliable, efficient, and timely delivery of data and processed products over the Internet is crucial to maximize the potential of ocean prediction systems for, among other applications, seasonal climate forecasting, ship routing, military planning, environmental studies, pollution tracking, and boundary conditions for higher-resolution coastal models.

The HYCOM consortium relies on OPeNDAP servers¹⁷ to distribute data and modeling products generated by the consortium. Currently, 60 TB of ocean-archived outputs, including currents, temperature, and salinity data from several simulations, and operational prediction systems, are readily accessible via the HYCOM OPeNDAP servers¹⁸. The use of these servers is growing rapidly, from 1 TB downloaded in 2007 to more than 4 TB downloaded in the first three months of 2009. It is important to keep in mind that these figures represent the subsets of data desired by the users. Had these users not been able to easily subset the data prior to acquiring them, the actual data downloaded would have been several magnitudes larger, impacting both the bandwidth to the archive and the users' computers; it is often difficult to manipulate, on the typical computer used by many researchers today, the large files in which satellite and model data are stored. An example of such a data set is the near-real-time output from the global ocean prediction system, which generates approximately 100 GB of data daily, including a four-day hindcast, an analysis, and a five-day forecast. The model output is converted to CF standard compliant netCDF files and made available via a DAP server typically within a day of the model run, allowing users to extract the subset of the global output relevant to their needs without having to download one or more of the *very* large files in which the daily output is stored.

¹⁷ <http://www.hycom.org>

¹⁸ <http://www.hycom.org/dataserver>

Asia-Pacific Data Research Center: Integrating Servers and Browser Clients

The Asia-Pacific Data Research Center (APDRC) was established within the International Pacific Research Center to increase understanding of climate variability in the Asia-Pacific region by developing the computational, data-management, and networking infrastructure necessary to make data resources readily accessible and usable by researchers, and by undertaking data-intensive research activities that will both advance knowledge and lead to improvements in data collection and preparation.

Most APDRC users are scientific researchers, and the data-serving system is based on the provision of direct, binary access to the data. OPeNDAP is used as the underlying system for this access. APDRC uses two OPeNDAP servers, one based on TDS and the other based on GDS. TDS is used to aggregate APDRC, netCDF-formatted data holdings. GDS, on the other hand, is used to serve data formats other than netCDF, including flat binary, GRIB, and HDF. OPeNDAP is also used as a back end to LAS and to DChart. These utilities provide users with Web-based browsing and simple plotting capabilities.

At present, APDRC serves about 80 data sets¹⁹ totaling approximately 20 TB. In 2007, more than 2 TB of data were downloaded via OPeNDAP.

Unidata

Unidata serves, via OPeNDAP, data from weather forecast model runs, National Weather Service radars, and some satellite-borne sensors. The model

data consists of over 10 collections, each with from four to 24 runs per day. Each day the archive receives approximately 100 GB of new data into a rolling, 30-day archive with, on average, 2 GB accessed via OPeNDAP per day.

Center for Ocean-Land-Atmosphere Studies

COLA serves the results of meteorological models via GDS. Typically, on a monthly basis, on the order of 1000 users make approximately one million data requests totaling in excess of 10 GB from COLA's servers.

LESSONS LEARNED WITH REGARD TO DISTRIBUTED DATA SYSTEMS

Taken together, the projects summarized in the previous sections provide a wealth of information related to accessing data in a distributed environment. In this section, we summarize this information as a set of lessons learned (or, in most cases, re-learned). These lessons are divided into two groups—those that were learned and addressed and those that raise issues that need to be addressed. The most significant lesson learned, from the collection of lessons summarized below, is that most of the obstacles to progress related to data access in a distributed environment are socially based rather than technically based.

Issues Identified and Addressed (If Necessary)

1. Modularity provides for flexibility.

The more modular the underlying infrastructure, the more flexible the system. This is particularly important

for network-based systems for which the technology, software, and hardware are changing rapidly. OPeNDAP was designed as a data access layer that lies above the data transport layer but beneath what is envisioned as a discipline-specific semantic layer. This structure allows the protocol to be used over a broad range of disciplines, providing the flexibility for uses other than those envisioned by the system designers. The interrelatedness of the examples presented in the previous section highlights the flexibility offered by the modularity of the approach adopted.

- 2. Data of interest will be stored in a variety of formats.* Regardless of how much one might want to define the format to be used by those participating in a data system, in the end, the data will be stored in the formats with which the users are the most comfortable. The same is true of use metadata. Furthermore, even when there is general agreement on the metadata convention, users will modify the conventions or will use them in unexpected ways.
- 3. It's tough to teach old dogs new tricks.* Introducing new technologies often requires a cultural change in usage that is difficult to effect. This lack of flexibility can have a negative impact on system development and system use. For example, many users would still prefer to save the data they acquire via OPeNDAP to a file prior to ingesting it into their application of choice, even when that application may be OPeNDAP-enabled.

¹⁹ http://apdrc.soest.hawaii.edu/cgi-bin/select_data3.s

4. *Some surprises encountered in the projects outlined above:*
 - a. The number of variables increases almost linearly with the number of data sets.
 - b. Users will take advantage of all of the flexibility offered by a system, sometimes to the disadvantage of all.
 - c. Even when using the same data format (e.g., netCDF), there is an incredible variability in the way the data are organized structurally. This lack of uniformity makes it difficult to write generic analysis programs to access the data sets.

Identified Issues That Have Not Been Addressed (Challenges for the Future)

1. *The structural representation of sequence (in situ) data sets is a major obstacle to interoperability.* The lack of a consistent organizational structure (as opposed to the format) for sequence data was the single largest constraint to the use of profile data in NVOODS.
2. *We know how to do it better.* Building a prototype distributed data system is straightforward because the developer either controls the “remote” sites in the prototype or has a close working relationship with prototype data providers. Problems arise when the system is opened to a broader class of data providers. The result is that new groups are continually developing the “next system,” which in turn results in a reduction in the adoption of existing systems and in confusion about the best path for the future. This issue has been the single biggest impediment to progress using OPeNDAP.
3. *The lack of a consistent structure for data inventories is a major obstacle to the use of distributed systems.* THREDDS provides a mechanism to address this issue, but until the community agrees on an inventory access protocol, this significant obstacle to data access via OPeNDAP remains.
4. *Time to maturity is order ten years, not three.* Developing new infrastructure takes time, both to iron out all of the details and for adoption of the infrastructure. This time frame is not consistent with typical federal funding cycles. A corollary is that *there is no well-defined funding structure for community-based operational systems*—it is much easier to obtain funding to develop a system than it is to obtain funding to maintain and evolve a system. This disconnect is a major obstacle to development of a stable cyberinfrastructure that meets the needs of the community, but this may be changing with the National Science Foundation’s Office of Cyberinfrastructure.
5. *The addition of server side operations, other than subsetting, is being approached by data providers in an ad hoc fashion.* Several groups (see section above on servers and Box 1) developed strategies for passing commands to the server that were not envisioned in the development of the original data access protocol. The fashion in which this sort of development is achieved and the subset of commands supported are not consistent from one group to another.
6. *There is disagreement on use versus discovery metadata.* The focus on metadata in the past has been on the use of these metadata to discover data sets as opposed to the use of metadata to facilitate actually making use of the data. The result is that there is much more agreement on *discovery* metadata than on *use* metadata. The netCDF community has addressed this problem through the development of the Cooperative Ocean-Atmosphere Research Data Standard and, more recently, the CF conventions. But, there are a number of other commonly used data formats for which there is less agreement on use metadata and even within the netCDF community there are many who do not use these metadata conventions. 

REFERENCES

- Cornillon, P., G. Flierl, J. Gallagher, and G. Milkowski. 1993. Report on the first workshop for the distributed oceanographic data system. Technical report, Graduate School of Oceanography, University of Rhode Island, 105 pp.
- Sirott, J., J. Callahan, and S. Hankin. 2001. Inside the Live Access Server. Paper presented at the 17th International Conference on Interactive Information and Processing Systems for Meteorology, January 14–18, 2001, Albuquerque, NM, American Meteorological Society.
- Sirott, J., D.W. Denbo, and W. Zhu. 2004. Dapper: An OPeNDAP server for in-situ data. Paper presented at the 20th International Conference on Interactive Information Processing Systems, January 11–15, 2004, Seattle, WA, American Meteorological Society.