

THE OFFICIAL MAGAZINE OF THE OCEANOGRAPHY SOCIETY

Oceanography

CITATION

Blower, J.D., F. Blanc, M. Clancy, P. Cornillon, C. Donlon, P. Hacker, K. Haines, S.C. Hankin, T. Loubrieu, S. Pouliquen, M. Price, T.F. Pugh, and A. Srinivasan. 2009. Serving GODAE sata and products to the ocean community. *Oceanography* 22(3):70–79, doi:10.5670/oceanog.2009.67.

COPYRIGHT


This article has been published in *Oceanography*, Volume 22, Number 3, a quarterly journal of The Oceanography Society. Copyright 2009 by The Oceanography Society. All rights reserved.

USAGE

Permission is granted to copy this article for use in teaching and research. Republication, systematic reproduction, or collective redistribution of any portion of this article by photocopy machine, reposting, or other means is permitted only with the approval of The Oceanography Society. Send all correspondence to: info@tos.org or The Oceanography Society, PO Box 1931, Rockville, MD 20849-1931, USA.

SERVING GODAE DATA AND PRODUCTS TO THE OCEAN COMMUNITY

BY JON D. BLOWER, FREDERIQUE BLANC,
MIKE CLANCY, PETER CORNILLON,
CRAIG DONLON, PETER HACKER, KEITH HAINES,
STEVE C. HANKIN, THOMAS LOUBRIEU,
SYLVIE POULIQUEN, MARTIN PRICE,
TIMOTHY F. PUGH, AND ASHWANTH SRINIVASAN



ABSTRACT. The Global Ocean Data Assimilation Experiment (GODAE [<http://www.godae.org>]) has spanned a decade of rapid technological development. The ever-increasing volume and diversity of oceanographic data produced by in situ instruments, remote-sensing platforms, and computer simulations have driven the development of a number of innovative technologies that are essential for connecting scientists with the data that they need. This paper gives an overview of the technologies that have been developed and applied in the course of GODAE, which now provide users of oceanographic data with the capability to discover, evaluate, visualize, download, and analyze data from all over the world. The key to this capability is the ability to reduce the inherent complexity of oceanographic data by providing a consistent, harmonized view of the various data products. The challenges of data serving have been addressed over the last 10 years through the cooperative skills and energies of many individuals.

INTRODUCTION

Oceanographic data are highly diverse, covering a broad range of spatial scales and encompassing remotely sensed data, in situ measurements, and numerical simulations. To extract the maximum amount of information from these data sources, it has been necessary to develop and deploy new technology platforms that allow data to be discovered, shared, visualized, and analyzed.

GODAE data holdings are very large: current ocean prediction systems and data repositories generate tens of terabytes of oceanographic data per year per organization. This rate is expected to increase rapidly with the deployment of new observing systems and increases in the resolution and complexity of numerical models. There is no single centralized authority with the resources or

mandate to oversee management of the comprehensive collection. Furthermore, transferring all of these data to a central location would be impractical and would force data providers to relinquish some control over their data products. Thus, a *distributed* data management approach that provides the ability to share oceanographic data effectively across the Internet is central to GODAE's aim of developing a global ocean forecasting system. In addition to solving many of the problems associated with managing large data holdings, a distributed system can provide users with more reliable and efficient data services through data replication and a more efficient use of networks.

The underlying approach to GODAE Data Services is a suite of tools based upon shared approaches. These tools

have been designed primarily for use by the scientific research community, but provide a solid foundation upon which other systems can be built to serve other user communities such as governments and commercial entities. The tools allow scientists to use data in a manner that frees them from the necessity to understand the low-level details of file formats, structure, or even the physical location of the data. The notion of *hiding complexity* is fundamental to the success of a data system that must deal with such large varieties and volumes of data. This hidden complexity must be balanced against the need to support a very wide range of applications, and it requires flexibility. GODAE services strike this balance by providing two different ways for the scientist to access and use data: (1) through Web portals, which hide data complexity but provide a fixed range of functionality, and (2) by allowing data to be ingested into the scientist's desktop tool of choice. These two approaches can work in concert, with the scientist performing preliminary discovery, evaluation, and analysis on the Web, then using specialist desktop tools, if required, to perform further tasks.

This paper summarizes the technological advances that have been made in the context of GODAE, advances that greatly facilitate the user's ability to discover, evaluate, visualize, download, and analyze a huge number of oceanographic data products.

GODAE DATA CENTERS AND THEIR PRODUCTS

The GODAE data centers (Blanc et al., 2008) provide a global framework of diverse data that have been produced by GODAE partners, including measurements acquired by remote-sensing satellites, observations from in situ instruments, and ocean analyses and forecasts produced by modeling and assimilation systems. These products have a wide range of spatio-temporal extents and resolutions to serve a range of possible applications.

Remotely sensed data are acquired globally and distributed in near-real time. They include data for monitoring sea level, sea surface temperature, surface wind, sea ice, ocean color, and, in the near future, surface salinity. These data are made available at several

levels of processing, ranging from raw data records to synthesized geophysical products such as ocean indicators and climatologies.

In situ data are acquired through various networks or programs, including the Argo network, the ship-of-opportunity program, and buoy networks. In situ data provide measurements of the ocean's interior state at depth, which is not observed via remote-sensing satellites. These data are all distributed in real time on the World Meteorological Organization's Global Telecommunication System (GTS), and are globally managed and distributed by Global Data Assembly Centers. Some centers, such as Coriolis (<http://www.coriolis.eu.org/>), also integrate data from different networks into coherent data sets for operational oceanography

purposes (assimilation, validation). Products may be individual measurement profiles or synthesized maps of global observations.

Ocean analysis and forecast centers produce data on global or regional scales, providing real-time forecasts (daily and seasonal) and historical time series. Numerical models resolve a variety of ocean features, from eddies to large-scale global circulation. Products may be hindcasts (simulations of past state), analyses (the best estimate of current state), or forecasts.

ELEMENTS OF THE GODAE DATA SYSTEMS

A Common Knowledge of the Products and Their Uses

It is important for all data providers and users to understand the processes that have been undertaken to produce a given data product. Knowledge of the entire context of a data product involves the traceability of events for operational production and the understanding of the product's attributes (including what is produced, the ocean region covered and its scale resolution, how it has been produced, who produced it, when and where it is made available, for how long, with what accuracy, delivery format and network delivery services, data and network service policy, and more). This information allows a user to decide upon a data product's fitness for a particular purpose. A key concept here is the notion of a product's "level," from raw instrument data (Level 0) all the way up to an ocean indicator (Level 5) (see Figure 1). The MERSEA Web site (<http://www.mersea.eu.org/>) uses this notion of "levels" to provide users with a consistent view of all available data products.

Jon D. Blower (j.d.blower@reading.ac.uk) is Technical Director, Reading e-Science Centre, Environmental Systems Science Centre, University of Reading, Reading, UK.

Frederique Blanc is responsible for ocean product information and dissemination for the Space Oceanography Division, CLS, Ramonville-Saint-Agne, France. **Mike Clancy** is Technical and Scientific Director of the US Navy's Fleet Numerical Meteorology and Oceanography Center, Monterey, California, USA. **Peter Cornillon** is Professor of Physical Oceanography, Graduate School of Oceanography, University of Rhode Island, Narragansett, RI, USA.

Craig Donlon is Principal Scientist for Oceans and Ice, European Space Agency/European Space Research and Technology Centre, Noordwijk, the Netherlands. **Peter Hacker** is Manager, Asia-Pacific Data Research Center, International Pacific Research Center, University of Hawaii, Honolulu, HI, USA. **Keith Haines** is Director, Reading e-Science Centre, Environmental Systems Science Centre, University of Reading, Reading, UK. **Steve C. Hankin** is Research Scientist, National Oceanic and Atmospheric Administration Pacific Marine Environmental Laboratory, Seattle WA, USA. **Thomas Loubrieu** is a research scientist at the Institut français de recherche pour l'exploitation de la mer (Ifremer), Plouzané, France. **Sylvie Pouliquen** is Project Manager, Ifremer, Plouzané, France. **Martin Price** is Data Scientist, Ocean Forecasting Research and Development, Met Office, Exeter, UK.

Timothy F. Pugh is Senior IT Officer, Centre for Australian Weather and Climate Research, Bureau of Meteorology, Melbourne, Australia. **Ashwanth Srinivasan** is Research Assistant Professor, Division of Meteorology and Oceanography, Rosenstiel School of Marine and Atmospheric Science, University of Miami, Miami, FL, USA.




Figure 1. From raw data to ocean summary information: various levels of ocean data products.

Base Technologies

Before sophisticated data systems can be built, there must be widespread adoption of common approaches for describing and transporting oceanographic data. Three main technologies play a large role in harmonizing GODAE products: the netCDF file format, the Climate and Forecast (CF) metadata conventions, and OPeNDAP (Open source Project for a Network Data Access Protocol).

File Formats and Conventions: netCDF and CF

A large number of file formats are available for expressing oceanographic data, from free-form, plain-text (ASCII) files, to highly structured binary formats. These formats often differ at a fundamental level, making it difficult to develop tools and applications that

work with all formats. To alleviate this difficulty, the GODAE community has standardized around netCDF (<http://www.unidata.ucar.edu/software/netcdf/>), which provides an array-oriented, platform-independent binary file format that can contain a wide variety of data types, from in situ measurements to large multidimensional grids of data from numerical models. The netCDF format is backed up by high-quality software libraries, in a variety of languages, which greatly ease the process of developing applications that consume and produce netCDF data. Furthermore, some of these software libraries (e.g., the official Java netCDF library) are able to read a variety of other file formats (such as GRIB—GRIdded Binary format, a World Meteorological Organization standard for encoding forecast data) and

interpret them as if they were netCDF files. In this way, the GODAE community has achieved harmonization of previously disparate data sets.

netCDF provides a simple, discipline-neutral data way to encode multidimensional arrays and their attributes. The CF conventions (<http://www.cfconventions.org>) provide the additional semantics defining how to encode oceanographic data (and data from other disciplines) in netCDF files. These conventions are currently focused on the description of gridded data from numerical models or analyzed satellite products. They provide a means to describe the grid on which the data are expressed, together with a suite of “standard names” that are used to identify the geophysical quantity that the data represent (e.g., “sea_water_potential_temperature”).

The netCDF file format and the CF conventions (known hereafter as CF-netCDF) provide an effective means of encoding oceanographic data. A number of tools for oceanographic data analysis and visualization have been developed on top of these technologies, including desktop tools (e.g., Ferret [<http://ferret.wrc.noaa.gov/>], CDAT [<http://cdat.sf.net/>], GrADS [<http://www.iges.org/grads/>], Ocean Data View [<http://odv.awi.de/>], and Ingrid [<http://iridl.ldeo.columbia.edu/dochelp/Documentation/>]) and Web-based tools (e.g., Live Access Server [see later discussion; Schweitzer et al., 2007], Godiva2 [Blower et al., 2009], and DChart [<http://www.epic.noaa.gov/epic/software/dchart/>]).

Although CF-netCDF has mainly been used to describe gridded data (e.g., from numerical models or satellites), use of CF-netCDF as a standard for encoding in situ measurements is rapidly gaining acceptance. Within the moored data (<http://www.oceansites.org/>), underway ship observation (<http://www.ifremer.fr/gosud/>, <http://samos.coaps.fsu.edu/html/>), and ocean profiler communities (Gould, 2005), standards that are based on CF are nearing completion. These standards build on CF by adding additional metadata needed to describe the specific measurements. The resulting files are fully CF-compliant and can be read by a number of generic CF-compliant applications. Note, though, that such files can describe only a single observation event such as a single time series, profile, or ship track. No widely agreed upon standards exist as yet to describe *collections* of observations (although a number of candidates exist), and this is a key obstacle to the

development of systems that allow users to visualize and process in situ data. The latest version of netCDF (version 4) contains new features that make it suitable for encoding such collections of observations, and research is ongoing into how this can be achieved in practice.

OPeNDAP

While netCDF provides a consistent data format in which to store GODAE data and CF provides a consistent metadata description of these data, an OPeNDAP service provides a consistent mechanism with which data may be accessed over the Internet (Cornillon et al., 2009). Specifically, users may access *subsets* of data sets residing elsewhere on the Internet and ingest them directly into their analysis packages. OPeNDAP servers may also provide for *aggregation* of large gridded data sets residing across several files. These capabilities are important, because a scientist wishing to access an oceanographic data set (for example, a multidecadal ocean reanalysis) often does not require the entire data set, which may be hundreds of gigabytes or even terabytes in size. Furthermore, it is often not desirable to regard the data set as a large set of individual files: the scientist may prefer to regard it as one large four-dimensional data set, which can be subsampled in numerous ways.

There is a very close relationship between the netCDF file format and OPeNDAP. It is possible to transmit netCDF data via OPeNDAP with (very nearly) no loss of information. Many desktop data analysis tools, such as Ferret, GrADS, and the MATLAB OPeNDAP Ocean Toolbox (<http://oceanographicdata.org/toolbox>), treat locally held netCDF data in exactly the

same way as remotely held data on an OPeNDAP server, providing the scientist with the capability to analyze and visualize huge quantities of distributed data. OPeNDAP servers can act as means of accessing data that are held by data centers in many other file formats such as HDF, GRIB, and BUFR. (Such formats are popular in other communities such as meteorology and Earth observation.) The end user does not need to know anything about which data format is used on the remote server. OPeNDAP is therefore a very powerful technology for data harmonization and integration.

Discovering Data

Each GODAE data provider has implemented a Web portal for users to discover and browse their products, and to provide users with links to download them. Dedicated catalogues exist at many of these sites to aid users in the discovery of specific data sets. This structure has led to the development of a large number of Web portals—each is designed differently, which is often confusing to users, particularly those outside the ocean community. Therefore, there are ongoing efforts to create integrated catalogues that provide users with a single point of discovery to an aggregation of data products held in GODAE archives. The MERSEA catalogue in Europe (by author Loubrieu and colleagues, submitted to the *Journal of Operational Oceanography*) is one example of such an aggregated catalogue.

Evaluating and Visualizing Data

Having found data of potential interest through a text search at one of the Web sites, the user will often like to evaluate these data for suitability for his or her

application before acquiring it. Many viewing services are now implemented, providing access to either predefined or dynamically generated visualizations. The low-level data standards described in the section on base technologies are extremely important here: it would not be feasible to provide visual access to all the diverse data sets in the GODAE systems without first agreeing upon how data are formatted.

Primary Viewing Services

In the European GODAE project MERSEA, the “primary viewing services” were defined to permit visualization of daily updated, predefined plots (historical plots may also be available). These services (see Figure 2) are simple to use, efficient, and fast. Images are pre-prepared based upon carefully preselected criteria, but such services provide limited or no capability for the

user to adjust the plots (e.g., to change the color scale or the region of interest).

Interactive Viewing Services

A recent advance in oceanographic data visualization is the development of interactive Web-based visualization systems. These systems allow the user to interact directly with the data and to customize the visualization to some extent (Figure 3). The Live Access Server,




Figure 2. Examples of primary visualization Web portals: BLUElink, FOAM, ECCO, and Mercator.

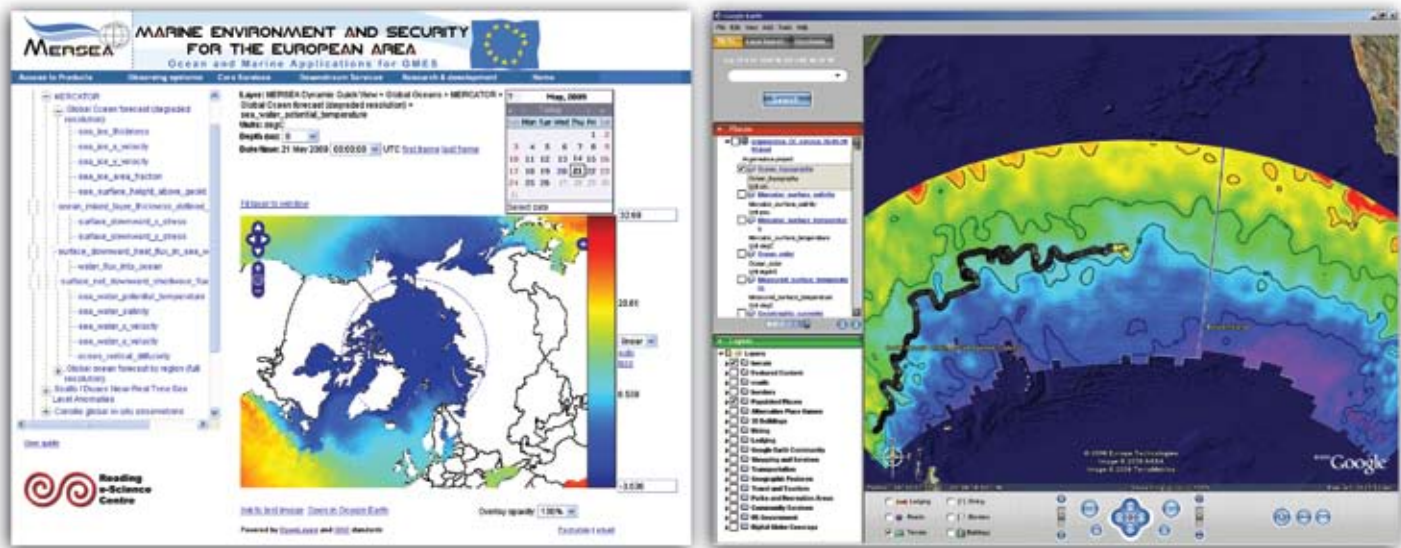


Figure 3. Interactive viewing services. Left: the MERSEA Dynamic Quick View system (<http://www.resc.reading.ac.uk/mersea>). This Web site allows the user to pan and zoom through large data sets, adjust the color palette and scale, change the map projection, and create animations. Right: Argonautica education project (<http://www.jason.oceanobs.com/html/argonautica>) demonstrating data visualization in Google Earth, which allows multiple data sources to be overlain.

discussed in further detail below, is a widely used example of such a system.

A new generation of Web-based visualization systems has recently emerged, based upon concepts and standards from the Open Geographic Information Systems (OpenGIS) community. This technology allows overlaying many sources of diverse information: for example, a user can overlay a forecast of sea surface height over a map of population density to gain a quick overview of hazards that may be posed by a forecast storm surge. This technology has been demonstrated in the European oceanography projects MERSEA (<http://www.resc.reading.ac.uk/mersea>) and ECOOP (<http://www.resc.reading.ac.uk/ecoop>), in which the Godiva2 system (Blower et al., 2009) has been adapted to create Dynamic Quick View (DQV) systems. This integration with GIS technology represents the beginnings of an important new direction in oceanographic

data analysis and visualization. It has the potential to allow oceanographic data to be combined with many other data sources from other communities, supporting a wide range of new applications in science and decision making.

Downloading Data

The visualization systems described above can help a new user to evaluate whether a data product meets his or her needs and to sift through the large amounts of available information. However, when the user knows which data products are required, he or she usually needs an efficient means of accessing the data. Two classes of technology are typically used: bulk file transfer and Web service.

Bulk file transfer systems are based upon FTP and/or HTTP to transfer unmodified files from the data center to the user. These systems can be secured by various means. Users with the

necessary access rights can be alerted to the presence of new data, which can then be downloaded. This simple download service provides an efficient and simple way to transfer whole data sets, file-by-file or pre-created derivative product files, either freely to all or through a registration process.

Data distribution by standard methods such as FTP or HTTP is reliable and mature, but can lead to a very high load on networks, as users need to download large data files even if they only wish to access a small subset of the data (e.g., to access data for a regional sea from a global ocean model). More sophisticated data services such as OPeNDAP allow much more powerful and flexible access to data, permitting users to access only the precise data they need. In addition to reducing the amount of data transferred, these “intelligent” data serving systems can help users by presenting data in consistent forms,

irrespective of the data format used by the data provider. The use of OPeNDAP services is not yet fully mature for operational use, but is very widely used by some user classes (e.g., research scientists). There is also much current development in the use of open geospatial Web services for delivering oceanographic data, with the particular aim of making data accessible to new communities (see earlier section on interactive viewing services). This technology is very new and presents many challenges; research in this field is ongoing.

In addition to the Internet-based access methods described above, there are also established forms of data provision, through dedicated communication channels such as GTS and Eumetcast. The technical differences between GTS and Internet-based services are profound: GTS is a “push” system in which data providers broadcast new data and consumers must sift all traffic for relevant information, whereas Internet-based services are “pull” systems in which data consumers choose which data they wish to download. Transitioning from one type of technology to the other presents many challenges.

Analyzing Data

Many dedicated software tools have been developed and made available over the past 10 years for various scientific applications (Blower et al., 2008). One of the primary goals of developing technologies for describing, discovering, visualizing, and accessing data is to allow data products from different providers to be inter-compared. In GODAE, some of the early inter-comparisons were implemented via the GrADS system, and this capability

was later implemented in Europe, the United States, and Australia using the Live Access Server.

The Live Access Server

The foundation of GODAE data management planning from its earliest days in 2002 included the use of OPeNDAP for access to distributed data sets and the use of the Live Access Server (LAS; <http://ferret.pmel.noaa.gov/LAS>; Schweitzer et al., 2007; Blower et al., 2008) to generate products (maps and scientific graphics, tables, and data subsets) and to perform intercomparison (regridding and differencing) between those data sets. Users interact with LAS primarily via a Web interface, behind which lies a data processing engine that can read data from many sources and process data in many ways. Having performed initial data visualization and processing on the Web, an LAS user can seamlessly switch to a desktop tool (such

fly to the user’s precise specifications of region, contour levels, scaling, color palettes, and other graphical styling issues. Latitude-longitude plots can be viewed on Google Earth (<http://earth.google.com>), from which the user can retrieve vertical profiles and time series plots by clicking on markers on the globe. LAS can also deliver subsets of data in various formats. The LAS animation viewer allows users to view maps, sections, and profiles as time animations.

The ability to intercompare fields—both by visual inspection and by computation of an anomaly field—is a fundamental requirement to support the GODAE effort. LAS offers two tools for this purpose: a “SlideSorter” for the intercomparison of fields within a single model output (or ensemble), and a regridding-and-differencing capability that can compare fields between model outputs hosted by separate institutions.

The complexity of analysis and visual-

“THE FUTURE SUCCESS OF OPERATIONAL OCEANOGRAPHY DEPENDS UPON THE COMMUNITY CONTINUING TO WORK TOGETHER TO AGREE UPON COMMON APPROACHES...”

as MATLAB or Ferret) for more complex and customized analysis tasks.

The most basic of LAS output products are custom scientific visualizations along all principal planes and axes: maps, time series, vertical sections and profiles, and Hofmuller (space-time) contour plots.

Figure 4 shows a montage of such outputs. All products are created on the

ization for Earth system model outputs has greatly increased with modern coordinate systems such as the tri-polar and cube-sphere grids. Recent versions of LAS address this complexity with “server-side analysis” capabilities that automatically perform complex tasks such as data regridding, thus simplifying user access to data.




Figure 4. A montage of selected output products from a Live Access Server.

CONCLUSIONS AND FUTURE OUTLOOK

GODAE operational oceanography capabilities have developed to the extent that, for global and regional seas, several centers now routinely produce data about the current ocean state, analyses, and forecasts. All partners have found that much of the operational infrastructure can be harmonized and shared to the benefit of all, moving from stand-alone to integrated multidisciplinary centers (e.g., MERSEA and the Global High-Resolution Sea Surface Temperature Pilot Project [GHRSSST-PP]); this has been one of the key lessons of the GODAE project. The future success of operational oceanography depends upon the

community continuing to work together to agree upon common approaches, in particular:

- **Common file formats and metadata standards.** The adoption of CF-netCDF has been very successful for the harmonization of gridded data such as numerical model output and satellite analysis products. The community now requires an equivalent standard for in situ and remotely sensed data; without such a standard, it will be very difficult to build the data systems and applications required by users of oceanographic data. Additionally, there needs to be wider agreement upon how to represent “high-level” metadata such as

the data production center, the spatial and temporal extent of data products, and links to further information such as documentation. Such metadata are very useful for discovery and attribution purposes, and for integration with GIS technology.


- **Integrated catalogues.** The profusion of different catalogues and Web portals for accessing oceanographic data can increase the difficulty to the user of discovering and accessing the data he or she needs. It will be very important for data providers to present their catalogue information in a form that can be harvested and aggregated by “meta-catalogues” that can act as single points of first

contact for new users. International standards are emerging that will help to enable the development of these meta-catalogues.

- **Common building blocks for applications.** This shared data-serving infrastructure will only flourish if it serves the ultimate goal of real user applications and commercial services. Consumers of oceanographic data often use specialized tools for visualizing and analyzing data. Several examples of GIS-based Web portals are already emerging, including Argonautica and AlgaeRisk. For these users, “general purpose” Web portals and desktop applications do not carry the specific functionality they require. The ocean community must develop good-quality reusable “building blocks” that can be assembled in various ways to enable new end-user applications to be developed at a reasonable cost. Such building blocks will include user interface components (such as interactive maps) and Web services for accessing data and catalogues.

There will be a continuing need to engage closely with users in order to ensure that future data systems meet their needs. The near future will see a strong move toward “operationalization” of the GODAE data systems. For example, in Europe, MyOcean, the Global Monitoring for Environment and Security Marine Core Service (Blanc, 2008), will deliver a pre-operational data system that will provide data, with specified service-level agreements, to users of oceanographic data in many disciplines all over Europe. It would not be possible to develop such a system without the advances made within GODAE.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge helpful reviews from Mike Bell, Ralph Rayner, and Gilbert Maudire. The Asia-Pacific Data-Research Center is funded under NOAA Award NA17RJ1230. 

REFERENCES

- Blanc, F. 2008. MyOcean information system. Paper presented at *EuroGOOS 2008*, May 20–22, 2008, Exeter, UK.
- Blanc, F., R. Clancy, P. Cornillon, C. Donlon, P. Hacker, K. Haines, S. Hankin, S. Pouliquen, M. Price, T. Pugh, and A. Srinivasan. 2008. Data and product serving, an overview of capabilities developed in 10 years. Paper presented at the *GODAE Final Symposium*, November 10–15, 2008, Nice, France. Available online at: <http://www.godae.org/Invited-papers.html> (accessed May 26, 2009).
- Blower, J.D., F. Blanc, P. Cornillon, S.C. Hankin, and T. Loubrieu. 2008. Underpinning technologies for oceanographic data sharing, visualization and analysis: Review and future outlook. Paper presented at the *GODAE Final Symposium*, November 10–15, 2008, Nice, France. Available online at: <http://www.godae.org/Invited-papers.html> (accessed May 26, 2009).
- Blower, J.D., K. Haines, A. Santokhee, and C.L. Liu. 2009. GODIVA2: Interactive visualization of environmental data on the Web. *Philosophical Transactions of the Royal Society A* 367:1,035–1,039.
- Cornillon, P., J. Adams, M.B. Blumenthal, E. Chassignet, E. Davis, S. Hankin, J. Kinter, R. Mendelssohn, J.T. Potemra, A. Srinivasan, and J. Sirott. 2009. NVIDS and the development of OPeNDAP. *Oceanography* 22(2):116–127.
- Gould, W.J. 2005. From Swallow floats to Argo: The development of neutrally buoyant floats. *Deep-Sea Research II* 52:529–543.
- Hankin, S., D.E. Harrison, J. Osborne, J. Davison, and K. O'Brien. 1996. A strategy and a tool, Ferret, for closely integrated visualization and analysis. *Journal of Visualization and Computer Animation* 7:149–157.
- Schweitzer, R.H., S. Hankin, A. Manke, J. Callahan, K. O'Brien, S. Du, and J. Li. 2007. Software architecture principles for Web-application infrastructure: How earth science software systems such as the Live Access Server fit in. Paper presented at the *23rd Conference on Interactive Information Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology*, January 14–18, 2007, San Antonio, Texas. Available online at: <http://ams.confex.com/ams/pdfpapers/118099.pdf> (accessed May 26, 2009).