CIRA

The Synergy of Science and Technology

Volume 20
Fall 2003

Colorado State University

Knowledge to Go Places

CIRA

Cooperative Institute for Research in the Atmosphere

Magazine Staff
Editor

Mary McInnis-Efaw

Design

Publications and Printing, Colorado State University

Technical Committee

Bernadette Connell, Doug Fox, Cliff Matsumoto, Donald Reinke

Fellowships in Atmospheric Science and Related Research

The Cooperative Institute for Research in the Atmosphere at Colorado State University (CIRA) offers a limited number of one-year Associate Fellowships to research scientists including those on sabbatical leave or recent Ph.D. recipients. Those receiving the awards will pursue their own research programs, collaborate with existing programs, and participate in Institute seminars and functions. Selection is based on the likelihood of an active exchange of ideas between the Fellows, the National Oceanic and Atmospheric Administration, Colorado State University, and CIRA scientists. Salary is negotiable based on experience, qualifications, and funding support. The program is open to scientists of all countries. Submitted applications should include a curriculum vitae, publications list, brief outline of the intended research, a statement of estimated research support needs, and names and addresses of three professional references.

CIRA is jointly sponsored by Colorado State University and the National Oceanic and Atmospheric Administration. Colorado State University is an equal opportunity employer and complies with all Federal and Colorado State laws, regulations, and executive orders regarding affirmative action requirements. In order to assist Colorado State University in meeting its affirmative action responsibilities, ethnic minorities, women and other protected class members are encouraged to apply and to so identify themselves. The office of Equal Opportunity is in Room 101, Student Services Building. Senior scientists and qualified scientists from foreign countries are encouraged to apply and to combine the CIRA stipend with support they receive from other sources. Applications for positions which begin January 1 are accepted until the prior October 31 and should be sent via **electronic** means only to: Professor Thomas H. Vonder Haar, Director CIRA, Colorado State University, humanresources@cira.colostate.edu. Research Fellowships are available in the areas of: **Air Quality, Cloud Physics, Mesoscale Studies and Forecasting, Satellite Applications, Climate Studies, Model Evaluation, Economic and Societal Aspects of Weather and Climate.** For more information visit www.cira.colostate.edu.

In Memory of Dr. Adrian Marroquin (1943-2003).

Adrian had been a CIRA Research Scientist working in the NOAA Forecast Systems Lab in Boulder since 1985. Prior to that, he was a CIRES Research Associate and a Postdoctoral Fellow at NCAR. His areas of expertise and experience included:

- *formulation and computer implementation of two- and three-dimensional numerical model dynamics for synoptic and mesoscale meteorological applications;*
- *computer implementation and verification of turbulence, convection, and radiation parameterizations and boundary layer formulations for numerical weather prediction models;*
- *dynamic meteorology and numerical analysis with applications to weather analysis and forecasting;*
- *teaching undergraduate physics and mathematics;*
- *teaching primary school students in his home country of Colombia.*

Adrian earned his B.S. in Mathematical Physics from the University Libre de Colombia, Bogotá, Colombia in September 1969, his M.S. in Physics from Florida Atlantic University, Boca Raton, Florida, in June 1973, and his Ph.D. in Atmospheric Sciences from the New Mexico Institute of Mining and Technology, Socorro, New Mexico in May 1981.

Our condolences and best wishes to Adrian's family, friends, and colleagues. He will be missed.

Technology at CIRA

By Ken Eis, Deputy Director

Progress in the atmospheric sciences is dependent upon a synergy of science and technology. Our understanding of the atmosphere includes theoretical physics and chemistry, but is practically limited by the quality of our observations and our ability to model and predict. Those limitations are more than likely due to technology: sensors too expensive to deploy in numbers, sensors with long-term drift and calibration problems, computers too small to model the physics and chemistry we already understand, and communications links too narrow to accommodate the flow of the data we can produce. Technology also hampers the understanding and distribution of weather information in the operational community.

This issue of the CIRA magazine is dedicated to the technologies that CIRA has been working on that support our research and subsequent operations. These technologies improve the quality and/or the efficiency of observing, analyzing, storing, and disseminating the scientific information and research that is the core of our mission. This magazine will highlight two forms of technology currently under development within CIRA.

Enabling Technologies

These are core technologies that improve either the infrastructure of the science support systems in CIRA, or provide a means to improve the dissemination, computation, or storing of information. Enabling technologies do not focus on what is being disseminated, computed, or stored, but rather generic technologies that can potentially affect many applications or functions. In some instances the technologies do not represent innovation in terms of function, but rather they are unique in that they represent improved efficiency or cost-saving ideas. Examples of these technologies to be discussed in this issue are:

Grid Computing – This technology uses existing computer resources (desktop administrative machines and scientific workstations not used at night) to perform large computational jobs. The only limitation is that the job needs to be partitionable. This technology offers a low-cost alternative to high-performance computing. Note that there are two articles on this technology.

DVD-Archive Systems – This technology is now offering a low cost method of storing the massive data sets produced by satellites. A year ago these systems cost tens of thousands of dollars and were only provided by turnkey "solutions" companies. Today they are made with commodity parts and can be built for less then $5,000.

Wavelet Compression Techniques – This technology has developed a new compression algorithm that allows both imagery and gridded data to be compressed for transmission. This is a unique application and unlike JPEG and JPEG 2000 (also Wavelet based) it allows the user to explicitly select the accuracy of the data.

Project Technologies

These are technologies that have been developed for specific scientific, educational, or operational purposes. Here the focus is on the end product or application and represents a specific use of technology.

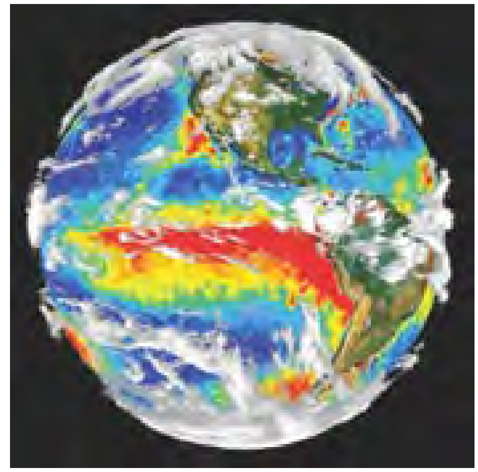

Examples include:

Volcanic Ash Coordination Tool Project – This system allows inter-agency coordination during the forecast process associated with volcanic events in the Alaska region.

Observation of Coral Bleaching Using IKONOS Satellite Data – This technology is exploring the use of very high resolution (better then 5-meter resolution) satellite data to determine the thermal heating stress on coral reef structures.

Contents

Technologies that Enhance the IMPROVE Network

By Christian Carrico, Jeff Collett, Derek Day, Doug Fox, Bill Malm, Taehyoung Lee, Xiao-Ying Yu, and Sonia Kreidenweis

The IMPROVE (Interagency Monitoring of Protected Visual Environments) network, as of June 2003, consists of 163 aerosol samplers, 40 nephelometers, 22 transmissometers, and 15 camera systems deployed to help characterize the ambient air quality and especially visibility of our National Parks, wilderness, and other sensitive locations. In addition to scientific applications, these data support regulatory programs mandated by the Clean Air Act including the relatively new regional haze regulations. The regional haze regulations require states to implement policies to reduce emissions so visibility at these pristine sites can return to natural conditions by the middle of this century. Additional details as well as a full reporting of the IMPROVE data can be found at the IMPROVE Website developed by CIRA (http//vista.cira.colostate.edu/improve).

The regional haze regulations require on-going determinations of how ambient visibility is different from "natural" visibility and the pollution sources that are impairing it. To help with these determinations, the IMPROVE network utilizes a set of aerosol samplers designed to monitor aerosol mass, size, and chemical constituents. Over the past 10 years these measurements have allowed determination of the major constituents of the visibility-impacting aerosol: sulfate, nitrates, organic carbon, elemental carbon, soil, and coarse mass. These data have been very useful in illustrating, for example, the dominant role played by sulfates in impacting the nation's visibility, especially in the Eastern US. Since sulfate predominantly comes from anthropogenic sulfur emitted from large, relatively easily identified sources, the network has proven value.

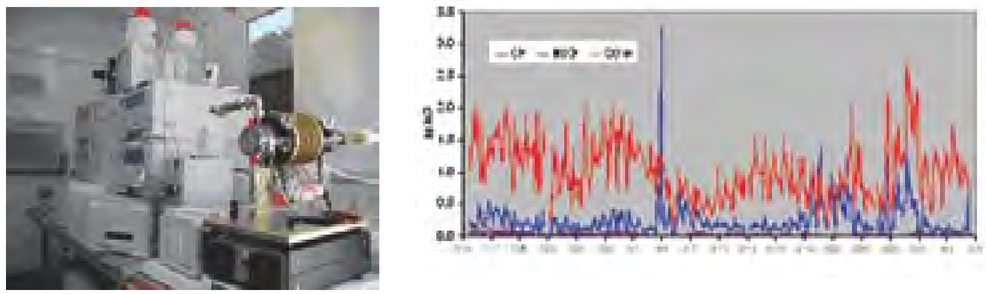

Figure 2 a, the PILS instrument and; b, an 8-week trace of 15 min PILS data collected at Yosemite National Park illustrating the anion components of the fine aerosol.

Following sulfates, nitrates and organics are the next most important aerosol components affecting visibility. In order to improve understanding of visibility impacts of these two aerosol constituents, the Department of Atmospheric Sciences at Colorado State University (ATS), CIRA, and others are conducting a series of studies to demonstrate new measurement technologies and their ability to quantify the roles nitrates and organics play on the ambient aerosol and visibility.

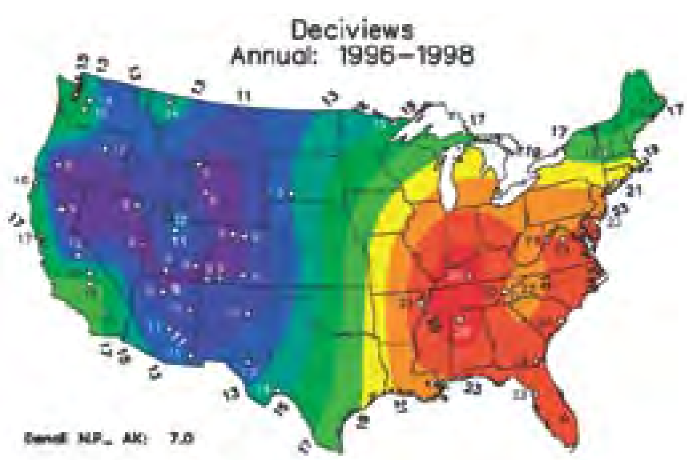

Figure 1, average visibility of the U.S. as determined from IMPROVE measurements.

Technologies

Two of the new technologies being investigated are high time resolution aerosol chemistry and hygroscopicity measurements.

High time resolution aerosol chemistry

Scientists from CSU/ATS's Atmospheric Chemistry group, including Prof. Jeff Collett, Dr. Xiao-Ying Yu, and Mr. Taehyoung Lee are making high time resolution measurements of aerosol composition. This involves using a **Particle into Liquid Sampler** (PILS) to measure water-soluble fine (smaller than 2.5 micrometers, known as PM2.5) aerosol ions. Aerosol particles are drawn into the PILS sampler where they are exposed to supersaturated water vapor, grow into droplets, and are impacted into a flowing liquid stream. This stream is fed directly to two ion chromatographs for measurements of aerosol anions (Cl-, NO2-, NO3-, and SO42-) and cations (Na+, K+, NH4+, Mg2+, and Ca2+). The PILS instrument is operated with a time resolution of 15 minutes, providing new insight into the temporal variability of aerosol ion concentrations. Figure 2 shows the PILS instrument and a sample of the measurements it is capable of making.

By examining how individual ion concentrations vary with time, it is possible to learn much more about the variability in aerosol composition (e.g., due to changing transport, interception of pollution source plumes, etc.). Measurements in the Midwestern U.S., for example, have helped identify rapid transitions between an aerosol dominated by ammoniated sulfate and an aerosol dominated by ammonium nitrate. PILS observations are also extremely helpful in understanding associations between different ions. For example, PILS data was key in revealing that nitrate measured in Yosemite National Park during summer 2002 was primarily present in the form of sodium nitrate, resulting from reaction of nitric acid or its precursors with sea

(continued on page 5)

IMPROVE Network *(continued from page 4)*

salt, and generally not present as ammonium nitrate. Yosemite PILS measurements of K+ ion, often used as a smoke tracer, also helped identify variability in smoke plume impact at the park. PILS measurements in Grand Canyon helped illustrate an association of nitrate (and some sulfate) with soil dust and sea salt aerosol while PILS measurements in southern California reveal a strong diurnal variability in aerosol ion concentrations and provide much better information about peak pollution levels.

Hygroscopicity measurements

The amount of water associated with the ambient aerosol has a significant impact on visibility. For example, the size of a pure ammonium sulfate particle grows dramatically when the ambient humidity is above 80%. This hygroscopic growth increases the scattering efficiency of the aerosol affecting visibility. Derek Day at CIRA is using nephelometers that control the amount of moisture the aerosol is subjected to in order to measure light scattering as a function of relative humidity to show how the bulk aerosol responds to changes in RH. Figure 3 shows the instrument.

The instrument consists of two nephelometers each with a separate inlet to control sample RH. One nephelometer always measures the "dry" (sample RH < 20%) scattering coefficient and serves as a reference signal. The second nephelometer measures the scattering coefficient at varying RH (20% < RH < 95%). The inlet of the drying/humidifying system consists of two sections in series; each section of the inlet is capable of modifying the sample aerosol's humidity separately. To begin a growth curve run, the aerosol is dried (generally about 15%) in the top part of the inlet then the aerosol is gradually humidified (generally up to about 90%) in the bottom part of the inlet. In this mode of operation it is possible to observe the aerosol deliquescence point. The system can also explore aerosol hysteresis by initially drying the aerosol in the top inlet and ramping the RH of the bottom inlet to about 90%. Then humidifying the aerosol in the top inlet to above 80% and subsequently reduce the humidity in the bottom inlet. Thus the aerosol sample is first exposed to high humidity then dried.

Figure 3, illustration of the specially configured Nephelometer for hygroscopicity measurements

Generally, there was reduced water uptake by ambient particles when the site was extremely hazy and being influenced by smoke from fires burning throughout the Western United States and there was significantly greater water uptake when the site experienced cleaner conditions and did not appear to be as impacted by smoke. Figure 5 shows a timeline of the scattering coefficient during the study (top plot) and the aerosol growth curve for each day (bottom plot). The range of sample RH for the aerosol growth curves in the bottom plot is from 10% < RH < 85%. Note that the highest scattering coefficients in the top plot generally correspond to the lowest f(RH) responses shown in the bottom plot. This suggests the aerosol associated with the smoke episodes is much less hygroscopic than the aerosol associated with the cleaner periods.

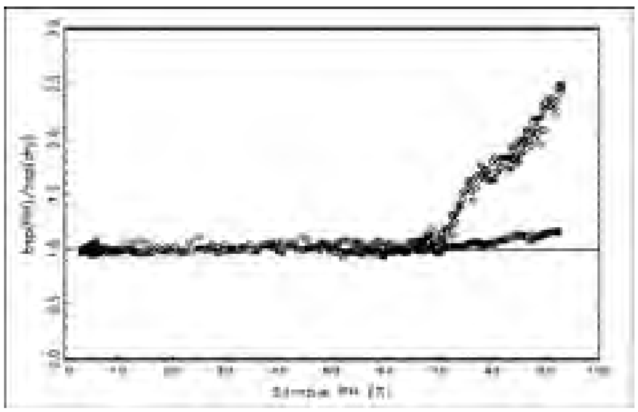

Figure 4 shows a wide range of aerosol growth curves that were observed during a study in Yosemite National Park in September 2002.

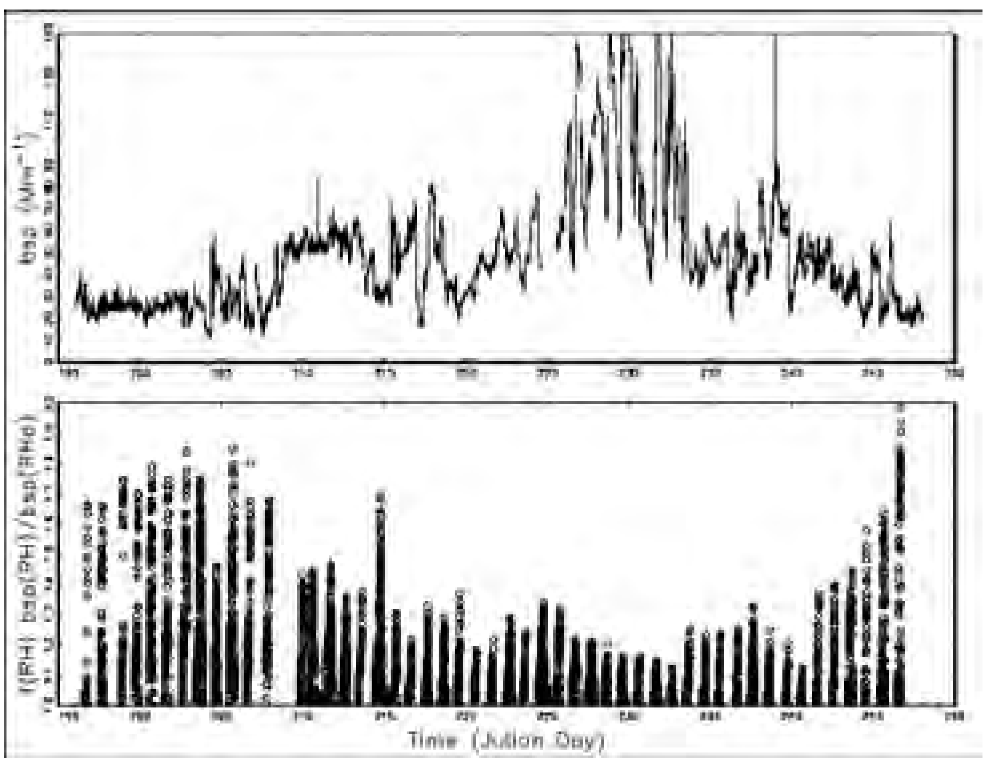

Figure 5, showing the light scattering coefficient (top) and f(RH) functions in time.

DVD Archive System

By Michael Hiatt

Since 1978, CIRA has collected Meteorological satellite data for research. The data is always archived for future use. Until recently, the only technology capable of storing the daily volume of data was tape.

In reviewing CIRA's archive, it was recognized that a new media was available called DVD. This media was developed to store home entertainment movies as a replacement for the aging VHS tape. In comparing the new DVD technology to the DLT tape technology CIRA was using, the following advantages were noted:

a. DVD media is one-third the price of the equivalent storage on DLT tape

b. A DVD writer is one-eighth the price of a DLT writer

c. A DVD reader is under $100.00

d. Archive verification is possible since the DVD can be accessed with no adverse effects unlike tape

e. DVD shelf life is expected to be at least three times better than tape

f. DVD's have no special environment requirements for storage

g. DVD data retrieval is random access rather than sequential

h. No special software is required to read a DVD

i. DVD's can be easily duplicated

These advantages clearly showed the need to develop a DVD archive system.

CIRA's Infrastructure group uses off-the-shelf technology to solve non-standard problems. This approach results in low cost, self maintainable solutions. Reliability is enhanced since the technology is in widespread use.

As a result, a DVD archive system was developed on a standard Windows 2000 based PC. This system supports up to six DVD drives and RAID-1 (two hard drives mirroring each other). The RAID system is an option to prevent the data waiting in the queue from being lost in the event of a hard drive failure. A custom software application, called DVDDC, was developed as the control software. Any number of these archive systems can be deployed depending on the data volume.

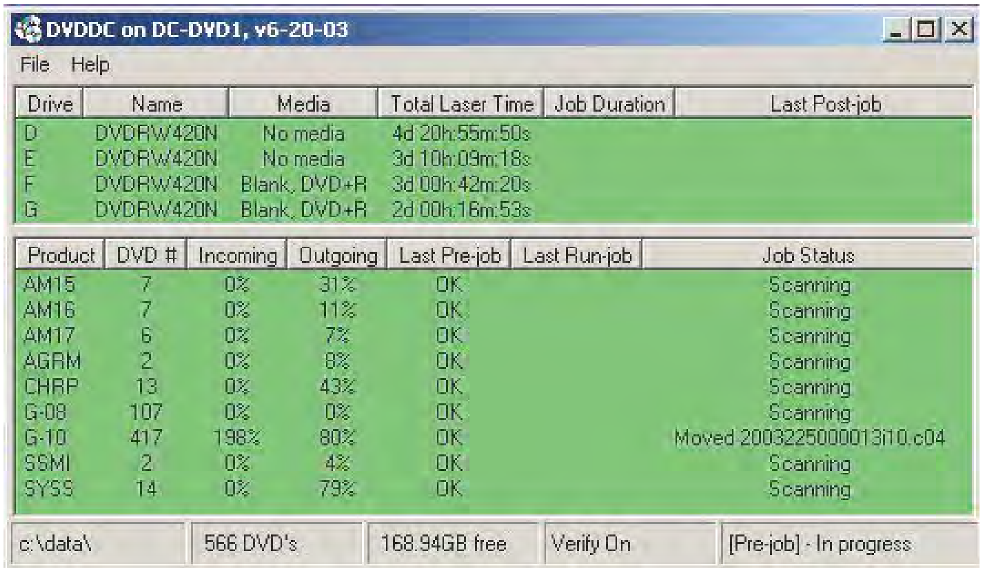

Figure 1, DVDDC Software

The DVDDC application, shown in Figure 1, is configured by the user defining any number of individual archives. Each of these archive's get a unique identifier and DVD starting index. The index provides an easy to use reference to the DVD's similar to a library catalog number. As data flows into the system, DVDDC sorts and queues the data. When the data volume reaches the user defined threshold, DVDDC automatically writes the data to the DVD. During the write process, several checks are made, including byte level verification. Regardless of the outcome, a log file is created and the DVD is ejected. If the DVD fails, a new DVD is created. If the DVD passes, the data is purged. A companion application keeps track of the log files and prints the appropriate labels on each DVD. The DVD's are then stored in binders for future use. The log files are HTML web files available on the Internet.

CIRA has been using this system since February 2003. During this time, more than 1,000 DVD's have been created. The archive staff reports that the DVD's are much easier to use and that the time required to restore data has been significantly reduced.

Figure 2, DVD archive system

Real-Time Three-Dimensional Lightning Display

By Phil McDonald

As part of the Range Standardization and Automation project, researchers at CIRA have been developing an operational user interface and data visualization display application for high-volume, real-time, three-dimensional LDAR (Lightning Detection and Ranging) data. Figure 1 illustrates the current display application that has been in use for several years at Cape Canaveral. This visualization includes a plan-view in the lower left panel, a view from the south in the upper left panel, a view from the west in the lower right panel, and a histogram in the upper right panel.

The new visualization, with its Tcl/Tk graphical user interface, is illustrated in Figure 2. While the layout of this new application maintains some similarities with the older application, there are several new features. The histogram has been moved to the user interface and the upper right panel now contains a three-dimensional rendering. The point of view of this three-dimensional view can be interactively manipulated by the user and is included to aid volume visualization.

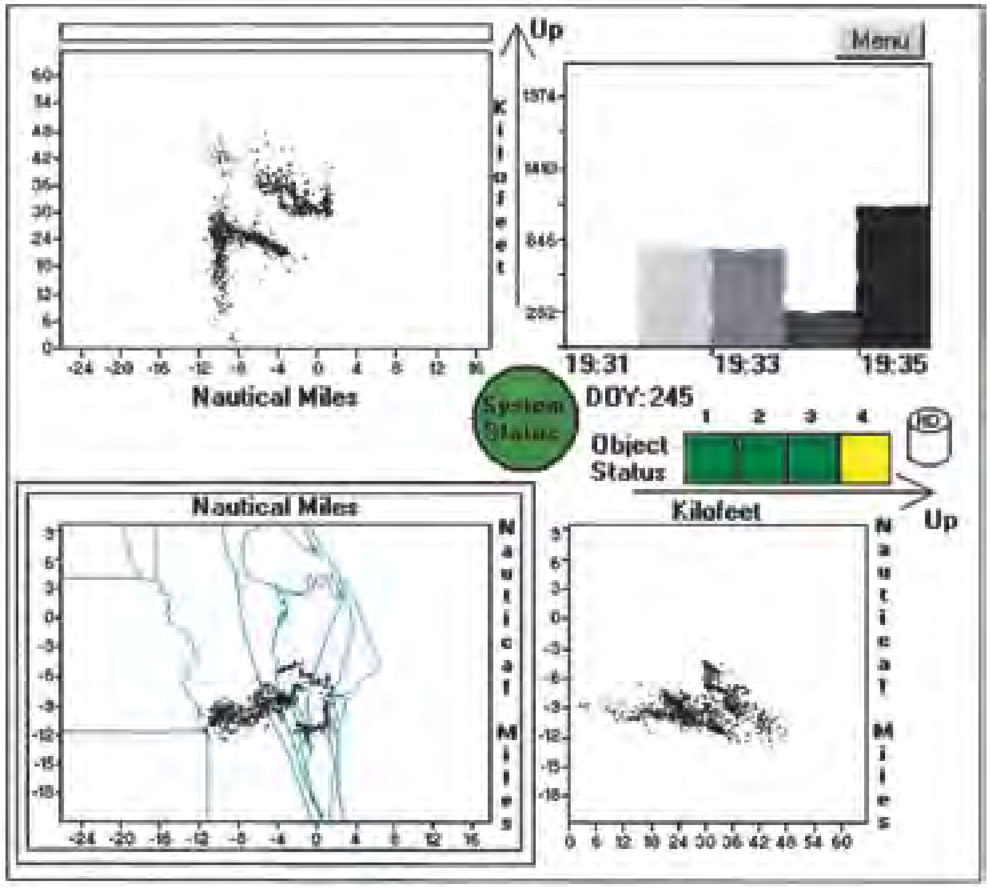

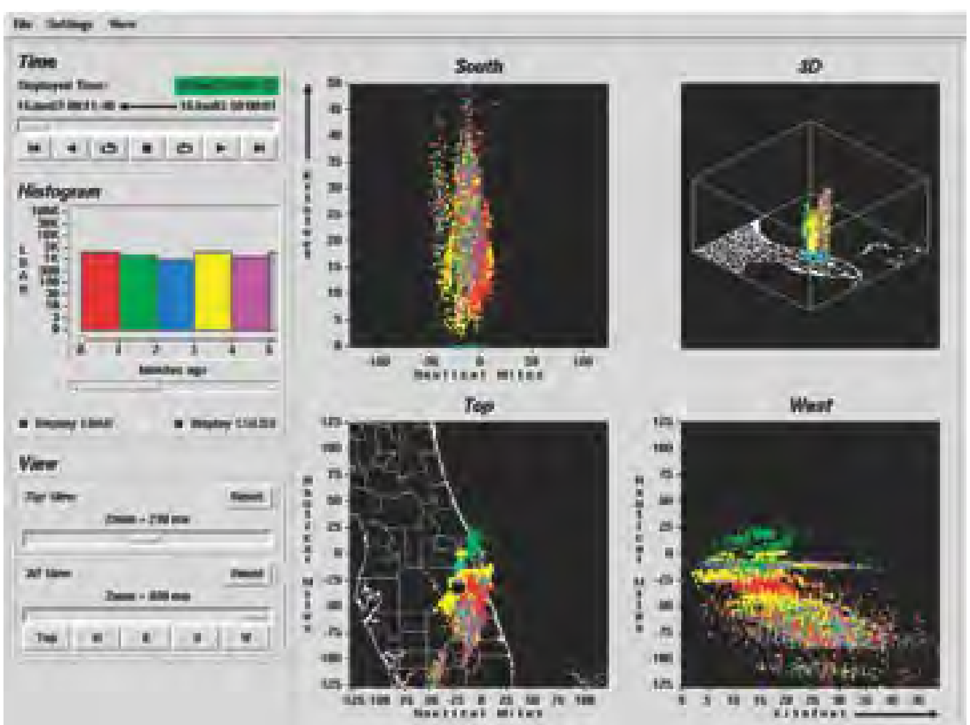

Figure 1, above. The current display application consists of a plan-view in the lower left panel; south view in the upper left; west view in the lower right; and a histogram in the upper right. Shades of gray indicate the age of the data: black for the most recent, light gray for the oldest.

Figure 2, left. Actual 3D and 2D data from sensors at the Kennedy Space Center from 16 June 2001. Small dots are plotted at each of the lightning step leader nodes and are colored according to their age: red 0-1 minute; green 1-2 minutes; blue 2-3 minutes; yellow 3-4 minutes; purple 4-5 minutes; gray more than 5 minutes. The histogram at the left indicates the frequencies of each of these age categories. Small cyan pluses ("+") are plotted at the locations of cloud-to-ground lightning strikes.

This new display application also uses color more effectively and has a higher degree of interactivity than the older application. To ensure optimum functionality, users of this application have been consulted on a number of occasions for their input.

Observation of Coral Bleaching Using IKONOS Satellite Data

By John Dietz

Coral reefs in many parts of the world have been subjected to a series of bleaching events over the past several decades. Bleaching occurs when coral polyps expel the pigmented symbiotic algae resulting in a bleached white appearance of living corals. The rate and prospects for recovery of the coral from bleaching is highly variable. Elevated water temperatures have been implicated in the majority of the large-scale bleaching events since the 1980s, and a number of SST indices have been developed to predict coral bleaching. One such product from NOAA is the satellite-derived Degree Heating Weeks (DHW) index, designed to indicate the accumulated thermal stress experienced by coral reefs.

A study was conducted at NOAA's National Geophysical Data Center to investigate using high spatial resolution IKONOS imagery (4-meter resolution in four spectral bands, 1-meter resolution in the panchromatic band) to detect coral bleaching. Lead NOAA investigators were Chris Elvidge and Alan Strong, and CIRA participants were research associates John Dietz and William Skirving. Because of the difficulty in acquiring timely field observations of bleaching, the validation of SST-based predictors such as the DHW has been limited. If the occurrence and severity of bleaching could be confirmed using high spatial resolution satellite data, this information could be used to validate the bleaching alerts issued and lead to improvements in the alert thresholds. Previous studies had cast doubt on the prospects of detecting and quantifying coral bleaching from space. This study examined the potential value of IKONOS imagery as a source of information on the frequency, severity, and extent of coral bleaching, as an aid to improving local and regional management and protection of reef resources.

An area of Australia's Great Barrier Reef (Keppel Islands) was chosen because of the availability of ground truth data and its high coral cover. The ground truth was data from underwater video transects conducted during the bleaching event, quantifying percentages and types of coral cover (see Figure 1).

The technique used for coral bleaching detection required comparison of a pair of IKONOS scenes – one pre-bleaching and one during the bleaching event. The best archive scene of pre-bleaching conditions was determined to be from August 2001, and a custom-tasked IKONOS collection was made during the bleaching event in April 2002.

Figure 1, video transect of bleached coral at Keppel Islands

Processing of the data involved co-registration of the imagery using image-to-image ground control points. Before the two images could be quantitatively compared, the digital number values (DNs) had to be brought into accord. A relative radiometric normalization was applied to the April (bleached) image using the "pseudo-invariant feature" (PIF) methodology. This procedure uses bright and dark pixel sets extracted from the reference and subject images to define a gain and offset for normalizing the subject image to match the radiometry of the reference image. These bright and dark pixel sets were used as input to a linear regression statistical analysis, yielding a gain and offset used to convert the raw April DNs into accord with the August DNs.

Difference images were generated for each spectral band using image subtraction. Also, average DN spectra were extracted for the areas of the video transects. When the spectra were plotted, the coral bleaching was detected as a brightening in band 1 (blue) and band 2 (green), but not in bands 3, 4, or the panchromatic band. By combining the first three image difference bands (1, 2, 3 as red, green, blue) to form a color composite and applying a contrast stretch, the location of the bleached coral could be readily discerned based on their gold color (see Figure 2).

This study confirms that the qualitative (presence/absence of bleaching) detection of coral bleaching is possible using IKONOS data, at least in areas of extensive bleaching. It may be possible to make detailed maps of living coral while the coral is in a bleached state, providing a useful capability of considerable interest to reef managers.

Figure 2, color composite image made from image difference data (R = Band 1, G = Band 2, B = Band 3) of the eastern portion of the Keppel Islands. Bleached areas show up as gold color.

Traffic Management Unit Project

By Jim Frimel

The TMU project is a Web-based research and development effort to develop products that are available directly to Air Traffic Controllers for their evaluation via the Internet. The major system used to acquire, distribute, create, and provide the required data sets for TMU is the AWIPS Linux data ingest and display system (Advanced Weather Interactive Processing System). Additionally, Web servers have been specifically developed for accessing and displaying these Web-based products.

The TMU project is currently in the initial phase of a four-phase effort designed to address the TMU's weather information needs relating to aviation hazards consisting of convection, icing, turbulence, and ceiling/visibility. Each phase will address the tactical (0-1 hour) and the strategic (2-6 hour) application of the above products to help the TMU decision maker in directing air traffic into and out of the ARTCC airspace. All phases of the research will follow an iterative process of defining, developing, demonstrating, and evaluating the weather-related hazard graphic and its presentation to Traffic Manager users.

Relating to advanced product displays, visualization, and the WWW, CIRA researchers have developed and are evaluating the prototype Tactical Convective Hazards Product (TCHP) on the Traffic Management Unit (TMU) restricted Website (Figure 1). The TCHP has been enhanced to allow toggle capabilities of map backgrounds and convective products based on feedback from the Traffic Managers use of the static TCHP displays. The convective SIGMET portion of the TCHP was enhanced to include convective SIGMET nowcast, convective SIGMET forecast, and convective SIGMET text. The convective SIGMET forecast is created by advecting the convective SIGMET nowcast using the motion information so that it is time matched with the National Convective Weather Forecast. A new impacted jet route map background was created using the convective products that comprise the TCHP for color-coding jet route segments (Figure 2). Testing, training, and evaluation plans have been created and added to the Website. The TCHP Viewer has also been enhanced to allow for looping.

The goal of the TCHP is to consolidate all tactical thunderstorm information into a single graphical product or limited suite of products for presentation to TMU decision-makers in an easily understood format. The TMU project will capitalize on development of advanced products from the AWRP and optimize the use of conventional advisories. Feedback from the ZFW Traffic Management Unit and Center Weather Service Unit participants will help refine the content and presentation. The Demonstration and Evaluation (D&E) will expedite fielding of advanced products by obtaining operational input early in the process. When there is agreement between the participants that a satisfactory product has been created, specific recommendations will be made for national implementation on FAA operational systems such as the Volpe National Transportation Systems Center Enhanced Traffic Management System.

For more information on these projects and others within the Development and Deployment Branch of the Forecast Systems Lab, please visit their Aviation Division's Home page: http:// www-ad.fsl.noaa.gov and select the link to the Development and Deployment Branch.

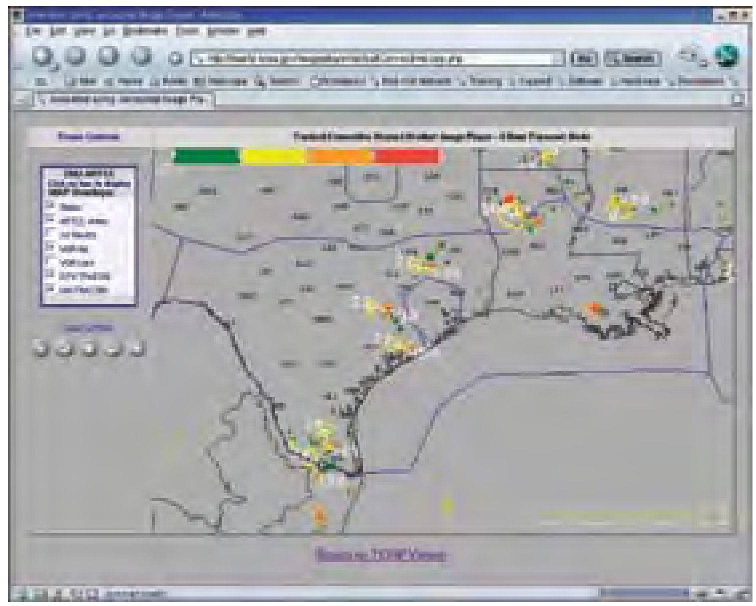

Figure 1 provides a view of the TMU Website TCHP Image Loop with default Convective Hazards.

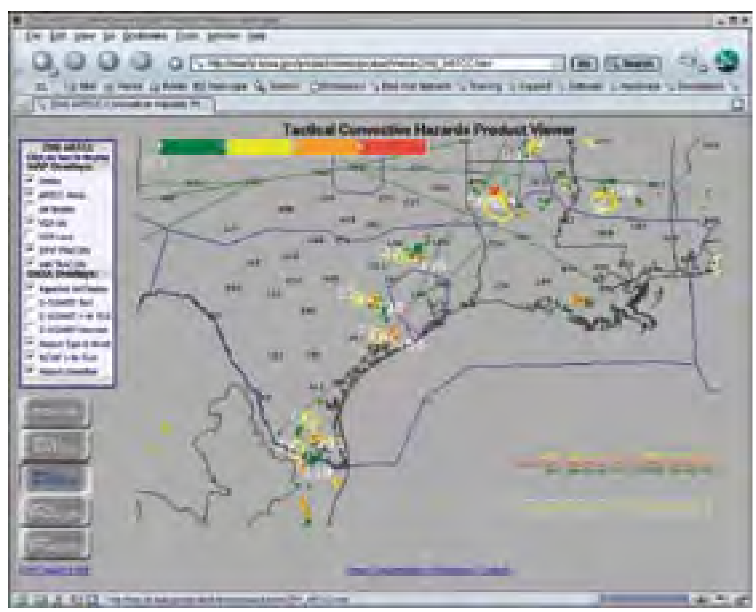

Figure 2 provides a view of the TMU Website TCHP showing the default ZHU ARTCC Scale with Impacted Jet Routes.

CloudSat Data Processing Center
Putting a Gorilla in a Shoebox

By Don Reinke

When CIRA was given the opportunity to be a part of the CloudSat team, we were presented with a significant technical challenge – to design, implement, and operate the entire data processing system for the life of the mission, all at a fraction of the normal cost of a NASA space mission. This article provides an overview of some of the innovations that will make that possible – and may provide some insights for other small research centers that are faced with the demands of managing ever-increasing volumes of satellite-derived science data and products.

The CloudSat Data Processing Center (DPC) is located in the new Atmospheric Science/CIRA Research Center (ACRC). The ACRC is collocated with CIRA and the Department of Atmospheric Science on the Foothills Campus of Colorado State University.

The CloudSat DPC will ingest approximately 26 GB of data per day and produce and store approximately 52 GB of products. In addition, the DPC will distribute a significant portion of those data to scientists and students around the world.

Perhaps the most challenging aspect of the DPC is the integration of science data processing software applications from a variety of locations. Nine groups of scientists, from four universities and a NASA lab, will write software to process CloudSat data into "Standard Data Products" (see table 1). All of these applications must run on the DPC system and interchange data, as many of the products depend on output from another product in the CloudSat data processing stream.

A second area that required some innovative attention is the merging of sensor data from multiple satellite platforms. CloudSat will be flying in formation with a number of other missions. This formation will provide a unique look at the atmosphere from different sensors. The DPC is tasked with managing the ingest and merging of these other data to allow the science teams to create products from the combined sensor data.

The last area of technical challenge is the automation of data processing, storage, and distribution. During the two-year CloudSat mission, two full-time people and a small staff of part-time students will operate the DPC.

The following is an overview of a number of technological innovations that we developed to handle these three challenges … but first a little background about the CloudSat mission.

CloudSat is a satellite experiment designed to measure the vertical structure of clouds from space and, for the first time, will simultaneously observe clouds and precipitation. The primary CloudSat instrument is a 94-GHz, nadir-pointing, Cloud Profiling Radar (CPR). The current launch date for CloudSat is March of 2005. CloudSat will be launched on the same vehicle with another ESSP mission called CALIPSO (a nadir-viewing, two-wavelength, Lidar).

CloudSat is a NASA Earth Sciences Systems Pathfinder (ESSP) mission. ESSP missions are relatively low cost missions that require the mission be built, tested, and launched in a short time interval – normally within three years. ESSP missions are peer-reviewed science investigations selected from proposals submitted to the NASA Office of Earth Science. They are led by a Principal Investigator (PI) who is often affiliated with a university. The PI for CloudSat is Dr. Graeme Stephens who is a professor in the Department of Atmospheric Science.

One of the key characteristics of ESSP missions is that they must fit under very tight cost constraints. Thus the mission segments, including the DPC, must develop reliable yet affordable solutions to each component of the mission. For a "normal" NASA mission, all of the data processing and distribution would be done by a major NASA center, such as those in Goddard or Langley. Our challenge was to simulate the operation of a major data center (the gorilla) on an "ESSP budget" (the shoebox).

First Challenge – Software Development at Remote Locations

CIRA is responsible for the implementation and testing of all of the software that is necessary to produce the nine standard data products shown in Table 1. The challenge, however, is that the software that creates each of these products is being developed by university or NASA JPL scientists at remote locations. Each software development team has their own favorite computer and programming language, and their own way of reading and writing the science data that is used by their software. In order to set up the CloudSat DPC under the time and cost constraints of the ESSP mission, CIRA and the Science Team Members have worked to together to develop a unique software implementation strategy. The plan works as follows:

1. All software development, including the system software at the DPC, is done on

(continued on page 11)

Standard Data Product	Description	Algorithm Dev.
1B-CPR	Level 1b CPR	JPL
2B-GEOPROF	CPR Geometrical Profile	U. Utah
2B-CLDCLASS	Cloud Classification	U. Maryland
2B-TAU	Cloud Optical Depth	CSU/Atmos
2B-LWC	Cloud Liquid Water Content	CSU/Atmos
2B-IWC	Cloud Ice Water Content	U. Utah
2B-FLXHR	Atmospheric Radiative Fluxes and Heating Rates	CSU/Atmos
2B-GEOPROF-LIDAR	CPR Geometrical Profile (CPR + Lidar)	U. Utah
2B-CLDCLASS-LIDAR	Cloud Classification (CPR + Lidar)	U. Alaska

Table 1, CloudSat Standard Data Products and responsible Algorithm Development Group

CloudSat Data Processing Center *(continued from page 10)*

the same platforms. Each system consists of an IBM PC, running the Microsoft Windows 2000 operating system, and a software package called "Developer Studio."

2. All of the standard data product applications are written in a common programming language. Each developer is using the FORTRAN 90 programming language, as implemented by Digital (Digital Visual FORTRAN).

3. A Web-based system, called the Algorithm Interface Management System (AIMS), was developed to manage the input and output specifications for each product (fig. 1). AIMS allows each software developer to specify the required input and output fields for their application. When an output field that is required by another application is modified, AIMS notifies the developer of that application that a change is pending to his or her input. The change is acknowledged by all those affected and it becomes a documented update to the output specifications for that product.

4. The main data processing system, called CloudSat Operational and Research Environment (CORE), does all of the required input and output of data fields for each standard data product application.

When the standard data product developer specifies in AIMS the data fields that are required as input, he or she can then click on a link in AIMS to create the necessary code modules that are used by CORE to ingest the required data and make it available to the application. The developer attaches a name that is used by their program to identify each input data field. When the main processing system reads in the required data, it attaches that name to the matching field.

In a similar manner, CORE creates the HDF-EOS data formatted output to match the output fields that the developer has specified in AIMS. This allows the developers to move very easily from software that runs on their independent test data sets to the standardized CloudSat data formats with a minimal amount of code modification.

Getting all of the developers to agree to this plan was a significant accomplishment. Many were used to working on different hardware/software systems and some had

Figure 1, screen dump of the 1B-CPR data product output fields as displayed by AIMS.

little expertise in the use of FORTRAN. However, in the short span of 18 months we have implemented all but two of the standard data products into the CORE system and the implementation has been extremely smooth.

Without the standardization of hardware and code, and the implementation of AIMS, the system would have taken at least twice as long (most likely, longer), and future modifications to software would have been very complex and costly.

A portion of AIMS is accessible to the general science community. You can access the application through the CloudSat DPC Website at http://cloudsat.cira.colostate.edu. You will be asked to enter a username and select a password on your first visit. Once you enter the site, select the "Data Products" tab and you will be able to access the output fields for each of the CloudSat science data products.

Second Challenge – Matching Data from Different Satellite Platforms

A unique aspect of the CloudSat mission is the fact that CloudSat will be flying in formation with other Earth Sciences missions. CloudSat will be a part of a constellation of satellites called the "A-Train" because they all follow the NASA Aqua satellite (fig. 2). The A-Train currently includes NASA's EOS Aqua and Aura satellites as well as a NASA-CNES (French space agency) lidar satellite (CALIPSO), and a CNES satellite carrying a polarimeter (PARASOL). This exciting new concept – formation flying – will add a level of complexity to an already challenging data processing task.

Most of the CloudSat standard data products are created, when possible, with data from other sensors within the A-Train system. To make these data available to the CloudSat algorithm developers, we implemented a software system that creates auxiliary data files. Each auxiliary data file (identified with a "-AUX" suffix in AIMS) contains data from another sensor that is geo-located with each CloudSat data point. The auxiliary data processor is called "generic" because it can take a wide range of data formats as input to create a matching -AUX file that can be used by any of the standard data product applications.

Examples of these formats (and currently available -AUX data files) are: HDF 4 and HDF-EOS 4 (i.e., MODIS radiance data, CERES, AMSU, etc.), and GRIB (i.e., ECMWF gridded data). This software application is extremely flexible and will allow the CloudSat data product developers to easily ingest data from each of the A-Train satellites (as well as other EOS mission data).

Third Challenge – Providing Data to Users Around the Globe

CIRA was also asked to take on the responsibility for distributing all of the CloudSat data to a wide range of users around the world. In order to handle this "Data Access and Support Center" (DASC) function, as defined by NASA, we developed

(continued on page 17)

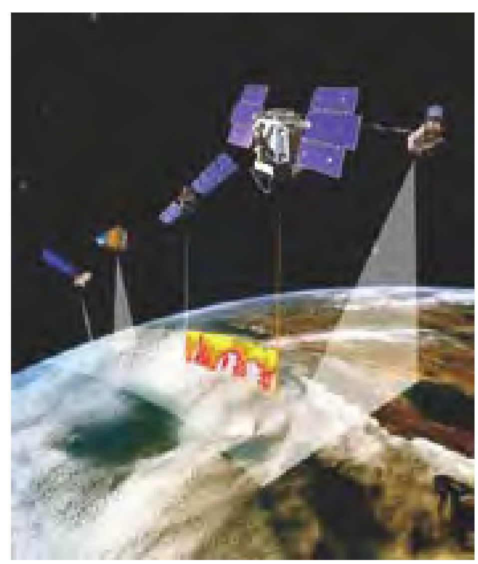

Figure 2, the Aqua "train" (A-train) is comprised of Aqua, as the leading satellite, followed by CloudSat (shown out of proportion), CALIPSO, Parasol, and Aura.

CIRA Communiqué: Employee News

2003 Research Initiative Award Winners Announced

In what has quickly become an important event every summer, the 2003 CIRA Research Initiative Award winners were announced at the Fort Collins "All-Hands" staff meeting July 31. This year two research groups were named, as well as an outstanding special award.

The VIEWS Project (Visibility Information Exchange Web System) is a group effort that is taking place on the National Park Service side of the CIRA house. **Rodger Ames, Shawn McClure, Shuxin Yin** and **Bret Schichtel** have combined their expertise to develop this multifaceted system. VIEWS consists of a database, a Website, supporting documents, and the procedures and policies necessary for the operation, maintenance, and continued evolution of the system as a whole. The objective of VIEWS is to assist states, Regional Planning Organizations (RPO), and the Environmental Protection Agency (EPA) with the increasingly sophisticated and complex analysis and interpretation of air quality data needed to support new regulations, including the Regional Haze Rule and new ambient standards for ozone and particulate matter less than 2.5 micrometers in diameter (PM2.5). The VIEWS system accumulates and maintains a database of specified air quality data and provides an Internet Website for the selection, presentation, analysis, and dissemination of these data. Starting as a project funded by the Western Governor's Association for only the Western Region Air Partnership (WRAP is the western US RPO), it was so well received that all of the RPOs agreed to adopt it and increase its scope and funding without competition.

The Research Award is acknowledgement of several factors. First, VIEWS started as a limited project (less than full time support for one person) and has grown to support for three persons. It has grown in scope, responsibility, and sponsorship directly as a result of the competence of the design and implementation team.

Second, the Web service is utilizing the latest technologies for its implementation. Application of new, state-of-the-art Web tools has been required in order to serve up the amount of data in response to complex random queries. This application is of such a level of complexity that one of the team is actually using the Web service for his Ph.D. dissertation in the CSU Department of Computer Sciences.

Thirdly, the Web service is actually innovating ways of working for states and other air quality scientists. Data analysis utilizing the tools and procedures developed and codified on the Web service are expanding and standardizing data analysis across the country.

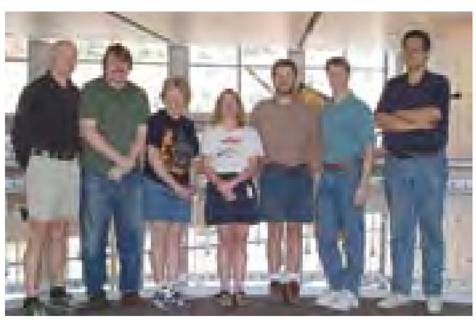

In our Boulder office, the Information and Technology Services (ITS) Data Systems Group comprised of **Chris MacDermaid** (Group lead), **Leslie Ewy, Paul Hamer, Bob Lipschutz, Glen Pankow, Richard Ryan,** and **Amenda Stanley** has also distinguished itself in the past year. They undertook many innovative projects to improve hardware/software performance, increase productivity, and decrease maintenance costs.

The enormous task of acquiring, processing, storing, and making available up to 2.4 TB of conventional (operational) and advanced (experimental) meteorological observations in real time is often taken for granted and goes unnoticed. The richness of this database, comprised of diverse observations from advanced automated aircraft, wind and temperature profiler, polar and geostationary satellites, GPS moisture, Doppler radar measurements, and hourly surface observations, is invaluable to researchers, systems developers, and forecasters within FSL and at various universities and other federal laboratories and agencies.

One sterling accomplishment of the DSG during the past year deserves special recognition – the design, development, and implementation of an advanced data handling capability based on state-of-the-art object-oriented programming technique. The new Object Data System (ODS) now makes it possible to process the rapidly increasing number of model grids and radar data in record time. Under the new ODS, one set of software handles all generic processing. No software needs to change when models change or when new models are introduced. The only maintenance needed is on the GRIB tables and the CDLs. Update to either of these does not require recompilation of software. Additionally, when a new model is introduced, ODS now generates the first draft CDL automatically. In many cases, this may be the CDL used for production. If not, it is changed because the researcher preferred a customized scaled-down version. Soon, ODS will permit researchers to use a Web interface to create a CDL and the netCDF files will be created from their customized CDL. The overall effort to implement the processing of a new model or changes to an existing model is now typically a few days – just a fraction of the time it took previously.

The overall effect is more responsive support to research projects, lower hardware costs, and greater flexibility due to the prevalence of open source tools available for Linux. In response to NOAA's call for quantifiable indicators of success, these much improved CIRA/DSG measures of perfor-

(continued on page 13)

mance are solid examples of our NOAA/FSL and CIRA collaborative partnership at work.

Finally, a special individual award was presented to **Cliff Matsumoto,** Associate Director in our Boulder office. Cliff is responsible for overseeing CIRA research conducted at NOAA's Forecast Systems Laboratory (FSL), at NOAA's Environmental Technology Laboratory (ETL), and at the National Geophysical Data Center (NGDC). There are approximately 50 CIRA employees working at the Boulder campus, a combination of research scientists, research associates, research coordinators, and state classified personnel conducting research in many different areas.

Among some of Cliff's outstanding contributions to the research teams in Boulder:

- Discussion and support of individual research projects: Cliff has a strong interest in the many research projects under his wing and he shows an insight into the details of the work.
- Bringing in new or replacing departing scientists as needed to ensure the continuation of on-going research projects: Cliff invests an enormous amount of time by participating in most of the selection and interview processes.
- Initiation of new research contacts: Cliff is a well-known figure in the meteorological community. In particular, his close relationship with meteorologists from the military has often resulted in new scientific collaborations.
- Interface CIRA – NOAA: Cliff's communication and management skills make him an effective link between CIRA and our federal partners.
- CIRA supervisor for a staff of 50 employees: In spite of his enormous workload, Cliff always makes his employees feel welcome to talk to him directly.
- Special committees: Cliff is a very active member of a number of CSU/CIRA committees. This service greatly benefits CIRA as an organization as well as individual employees.
- Publications / Reports: Cliff's scientific background, his writing skills, and his attention to detail are of great value when it comes to writing papers, etc.
- GLOBE CAN Proposal: Cliff was a key contributor to the GLOBE CAN proposal. As a result, the GLOBE CAN was awarded to the CIRA/CSU/UCAR team in the face of strong competition. This award made CIRA and CSU an important systems, education, and outreach partner in the new GLOBE Program.

In short, Cliff is an exceptional leader, manager, and researcher. He supports a wide spectrum of cutting-edge research activities, and his management and communication style results in an extremely pleasant working atmosphere. Cliff's noteworthy accomplishments have had a substantial impact on the research mission at CIRA.

NOAA Research Team Member Employee of the Month – April

Brent Shaw, a CIRA employee based at the NOAA Forecast Systems Lab in Boulder, was selected as Employee of the Month for April 2003. This marks the second consecutive month that a CIRA employee was honored by our NOAA federal partners! Brent was presented this award for his role in developing a numerical weather prediction system for the Space Centers at Vandenberg AFB, California and Cape Canaveral, Florida. This system enhances weather forecasting for space operations, particularly at launch time when low-level winds and temperatures are critical to operations.

As the official announcement read: "Mr. Shaw led an effort to establish a unique, high-resolution, local numerical prediction capability at each facility that allows diagnosed clouds and precipitation in the model's initial condition, leading to large improvements in short-range forecasts of clouds and precipitation. He merged the output with the AWIPS LINUX-based workstation at the facility, to ensure the operational staff has ready access to these important forecast products. This effort required expertise in many technical matters, including hardware configuration, knowledge of data structures, model code, operational scripts, post processing, and product generation. Mr. Shaw completed all this work in a superior manner, and optimized the model configuration to exceed the Air Force's requirements within computer resource constraints."

Recent Promotions at CIRA

CIRA management was pleased to recognize the outstanding accomplishments of 14 employees with special promotions this past July 1. These promotions are a part of the new career ladder recently implemented by Colorado State University for Administrative Professionals. Prior to the ladder, there was no system in place for APs to advance; however, the concerted efforts of CSU administration in conjunction with input from the Administrative Professional Council was successful in the establishment of the Research Associate/Research Scientist ladder.

In the CIRA/NPS division, **Mike Barna** was promoted to Research Scientist II recognizing his research accomplishments developing and applying regional air quality modeling tools. **Derek Day** was promoted to Research Associate III based on his contributions to aerosol measurement instrumentation development and application in various field programs.

Shripad Deo was promoted to Research Associate III in recognition of his work with the National Weather Service River Forecast Office in Kansas City, KS.

(continued on page 20)

Data Assimilation: Why and How?

By Tomislava Vukicevic

Measurements of atmospheric conditions are performed regularly around the globe several times a day via networks of ground-based weather stations, aircraft, and remote sensing platforms. The total number of weather data from all the networks is large. For example, the National Center for Environmental Prediction (NCEP) collects about 10^6 observation data points every three hours. This number of observations is not, however, sufficient to characterize the entire atmospheric state since this represents an observation for every 197 square miles.

The atmospheric state is very complex as it is a collection of air, water vapor, and a variety of particles characterized with continuous fields of energy and mass in motion. Temperature, pressure, and concentration of water vapor and particles including other phases of water in the atmosphere is varying in space on scales as small as a millimeter. This suggests that a full description of the atmospheric condition at any time would require as many pieces of information as there are millimeter unit volumes in the atmosphere. In addition, to know how this very large state changes with time it would have to be recorded very frequently, for example every few seconds or minutes. Thus, the total number of pieces of weather information for the complete description of the atmospheric state is apparently orders of magnitude larger than the number of available weather measurements.

Do we need to know it all?

Is the information of atmospheric state on millimeter spatial scales and second time scales necessary for good weather analysis and prediction? The answer to this question depends on what we hope to achieve. Weather phenomena are characterized with a spectrum of spatial scales from millimeters to the entire globe. Each of these scales has specific requirements for a sufficient number of measurements to fully describe it. For example, synoptic-scale weather phenomena such as cyclones and anticyclones and the associated continental scale fronts could be well represented with measurements having horizontal spatial resolution of about 100 km. Also, the synoptic-scale measurements should include vertical scales of about 100 m in the lower troposphere to 1000 m in the higher troposphere. This together requires a volume of about 10^7 measurements assuming they are readily available with approximately even distribution in space, and that all physical quantities needed to describe the atmospheric state on synoptic scales are measured at the same time. This estimated number is compatible with the current number of weather measurements but these are not evenly distributed in space, nor are all quantities measured at the same time. Many are measured directly. In fact more than 99% of the weather measurements today are provided by satellite remote sensing, which implies that they do not represent direct measurements of the atmospheric state such as temperature, pressure, and humidity.

This implies that some procedure must be applied to the measurements to map them into a regular distribution in space, referenced at the same point in time from which the weather prediction is then produced for the given range of scales. This procedure is known as the **weather data analysis**. When the weather data analysis is constrained by known governing equations similar to a numerical weather prediction models, it is called **data assimilation.** Thus, data assimilation assumes the utilization of a physically-based model of the atmosphere in producing weather analysis.

The data assimilation procedure is not only required to produce weather data distributed on a regular spatial grid, but must also include objective criteria for maximizing the accuracy of the resulting analysis with respect to uncertainties in the measurements and the assumed model of the atmosphere. The measurements and model are characterized with imperfections (i.e., the errors) but if combined in a correct way they could produce the weather data with errors smaller than the input uncertainties. This improvement is possible, however, only if the number of independent weather measurements is larger than the number of independent data points in the modeled atmosphere. This is hardly achievable even in the example of the synoptic scale weather phenomena. The problem of insufficient number of measurements in the data assimilation is solved by introducing what is known as **prior information.** This information represents the weather analysis that has been produced **prior** to the assimilation of new measurements. The prior information is also characterized with errors which must be known. The knowledge of prior errors is acquired in practice by repeated experience of the data assimilation procedure.

In the current weather data assimilation practices, the ratio between the number of actual independent measurements and the model grid points rapidly declines with decreasing spatial scales of interest. For example, to produce weather analysis and prediction for severe weather, the typical spatial scales in the model should be on the order of hundreds of meters to several kilometers, i.e. mesoscale. Thus, for the mesoscale, the new observations not typically used in the synoptic scale weather analysis should be considered as well as different prior information.

There are few types of weather measurements that provide the mesoscale spatial resolution over extended regions. They are typically in the domain of satellite remote sensing. Remote sensing measurements have been used in modern weather data assimilation in the last 10 years, but mostly in synoptic scale weather analysis. The use of satellite remote sensing for the mesoscale data assimilation is a new area of research in the atmospheric sciences. Satellite remote sensing measurements are rapidly improving in the spatial and spectral resolution, implying great potential for mesoscale weather data assimilation.

Data assimilation research in CIRA

In CIRA the data assimilation methods are investigated that specifically include high resolution weather models and a variety of satellite and other measurements about the meso and synoptic scale weather processes. Large portions of this research are sponsored by the Department of Defense through the Center for Geosciences at CIRA, but also by research grants from the National Science

(continued on page 20)

Transmission of Current GOES Imagery aboard the NOAA WP-3 Aircraft Using Wavelet Compression Techniques

By John Knaff, Ning Wang - CIRA; and Mark DeMaria, Raymond Zehr - NOAA/NESDIS

NOAA's two WP-3 aircraft, Fig. 1, are the primary tools for the annual hurricane field program conducted by the NOAA Hurricane Research Division (HRD). The WP-3 aircraft are also used for operational reconnaissance missions to supplement the flights of the U.S. Air Force Reserve, which operates out of Keesler Air Force Base in Mississippi. During the hurricane off-season, the WP-3 aircraft are used for many other atmospheric research missions throughout the world. The WP-3 aircraft are instrumented to collect flight-level atmospheric data, can release dropwindsondes to obtain vertical profiles of atmospheric parameters, and have on-board Doppler radars. However, the WP-3s do not have the capability to display or animate current satellite imagery because the aircraft have limited bandwidth communication systems. This capability would significantly aid many of the WP-3 research and operational missions.

Because of the potential beneficial impacts of satellite data, a demonstration project was undertaken during the 2002 and 2003 hurricane seasons to display and animate real-time GOES satellite data aboard these aircraft. This project combines the efforts currently underway at HRD to increase the communication capabilities of the WP-3s using cell-phone and Internet technology, and recent advances in data compression technology being developed at NOAA's Forecast Systems Laboratory (FSL). Even with the recent WP-3 to ground communication improvements, the bandwidth is still very limited (maximum rate of 2400 baud), although it is a dramatic improvement over what was previously available. To accomplish these goals, advanced wavelet data compression techniques are applied to GOES satellite imagery that allow the transmission of large data volumes associated with satellite observations over the very limited bandwidths available for communication with the aircraft.

Figure 1, one of the two NOAA WP-3 Aircraft.

Aircraft to Ground Communications

The communications to and from the WP-3 aircraft and the ground is currently provided through INMARSAT satellite communications (SATCOM), which provides voice communications and data transmission rates of 2400 baud. This bandwidth enables the transmission of only limited amounts of real-time data. For instance, during the NOAA's 2002 Hurricane Field Experiment the SATCOM demonstrated the capability to transmit a 42 kilobyte (kb) Joint Photographic Experts Group (JPEG) satellite image in approximately two and a half minutes (see Figure 2). While the current data transmission rates would easily enable the transmission of half-hourly GOES data to the WP-3 aircraft, faster data transmission is still desired.

Data Compression Techniques

The choice of compression technique is mostly based upon a tradeoff between compression ratio and image quality. For the purpose of transmitting satellite imagery it is very important to have a technique that creates imagery that has a high degree of fidelity without shadows and distortions and

Figure 2, a laptop displaying the 42 kb water vapor JPEG image that was downloaded during a WP-3 flight into Tropical Storm Hanna on September 13, 2002. It took approximately 2 minutes to download the image. Also shown is Mike Black of HRD and Sean McMillan, a member of the Aircraft Operations Center NOAA-43 flight crew.

contains as much useful information as possible. For example, the JPEG format is a commonly available image compression technique. While JPEG compression achieves the compression needed for this project, the fidelity of the decoded image is far from satisfactory for this application. JPEG-compressed images usually exhibit some noticeable blocky effects at a compression ratio as low as 10:1.

A similar compression problem arose during the development of FX-Net (Wang and Madine 1998) at FSL. FX-Net is an Internet-based meteorological workstation with an AWIPS-like user interface, which operates with only modest communications and computing capability. After testing a variety of compression methods, the wavelet transform was chosen as the approach to image compression for FX-Net. The wavelet transform was introduced in the early 1990s and has remained a cutting-edge technology in image compression research (Daubechies 1992, Prasad 1997). Like the Fourier transform, the wavelet transform relies on a particular set of basis functions. However, the set of basis functions that the wavelet

(continued on page 16)

Transmission of Current GOES Imagery *(continued from page 15)*

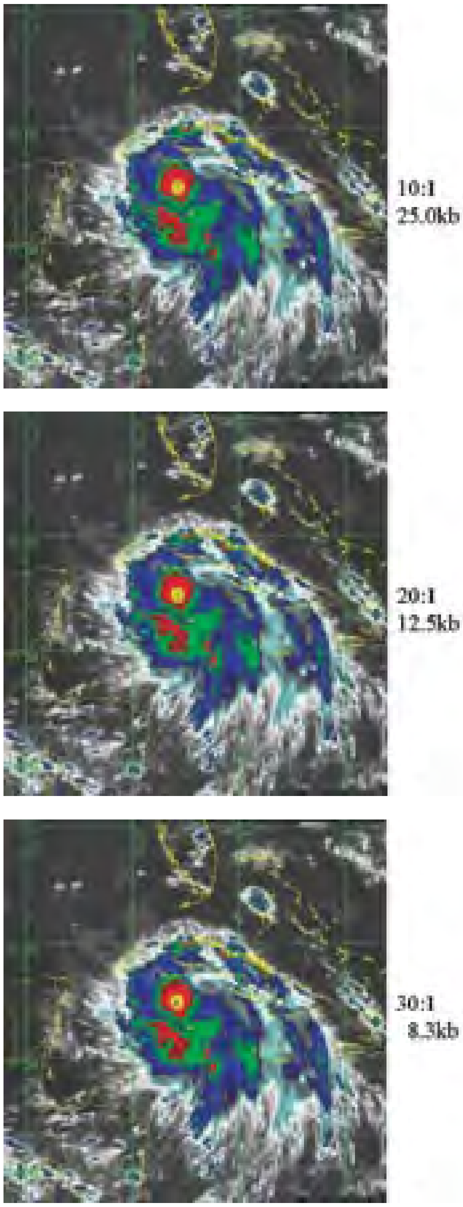

Figure 3, examples of images resulting from wavelet compression using compression ratios of 10: 1,25kb (top), 20:1,12.5kb (middle), and 30:1, 8.3kb (bottom).

transform uses is localized in both space and frequency, whereas a Fourier transform only contains frequency information. The ability of the wavelet transform to contain some spatial information, in addition to the frequency information, allows it to achieve excellent compression of meteorological images.

The wavelet data compression usually consists of three steps: transform, quantization, and entropy encoding. The transform step is done through a fast discrete wavelet transform algorithm. The quantization step finds the "best" representation of the image for the specified compression ratio by minimizing the mean square error in the coefficient domain. The last step, entropy encoding, further reduces the size of the compressed data set, in lossless fashion.

Examples of the images resulting from wavelet image compression routines with various compression ratios are shown in Figure 3. For comparison the GIF and JPEG (quality = 30) renditions of the same image (date time) are shown in Figure 4. The images decoded from the wavelet compression show an ever-decreasing degree of higher frequency detail with increasing compression ratio. However, the resulting images are still superior to the JPEG image, which actually has a larger file size.

A Demonstration Project in 2002 and 2003

To demonstrate the ability to transmit and animate satellite imagery aboard the WP-3 aircraft, tropical RAMSDIS (Molenar et al. 2000) was utilized to automate the process of acquiring, compressing, and serving the satellite data on the ground, a workstation on the WP-3 acquired the compressed files (via ftp), decompressing them, and then loops the resulting imagery in a Web browser.

With the purpose of improving reconnaissance and data collection capabilities aboard the WP-3 aircraft, only a fraction of currently available satellite data is utilized. Three spectral channels of the Geostationary Operational Environmental Satellite (GOES) were acquired half-hourly from the NESDIS severs for this purpose. Visible, 11 µm (IR4), and 6.7 µm (water vapor) imagery were remapped to Mercator projections at 2 km, 4 km, and 8 km resolutions, respectively, resulting in three 500 x 500 pixel, 1-byte datasets with navigation, calibration, and time information with a size of (251 kb).

From these raw data files, the 1-byte image information was extracted to create a data file (225 kb) and from the navigation and time information, a minimum set of information was extracted and put in a navigation file (52 bytes). The navigation file contains the number of points in the x-direction, the number of points in the y-direction, upper-left latitude, upper-left longitude, lower-right latitude, lower-right longitude, base latitude for the Mercator projection, Julian Day, UTC hour, minute, and second. These two files along with the compression ratio provided input to the wavelet compression algorithms that output a file containing the encoded wavelet transform information and the navigation information (WLT). This file is then made available to the plane via ftp.

A script on the WP-3's computers then contacts the ftp site to look for new WLT files. If newly created WLT files exist, they are transmitted to the plane and decoded into GIF satellite images or raw data files. Each of the three image products is then sorted with respect to time and made available to three separate HTML-based JavaScript programs that loop the image products in almost any HTML browser.

Examples of what the browser loops of the three satellite products looked like for Hurricane Isidore as it was approaching the Louisiana coast on 25 September are shown in Figure 5. Currently, four hours of half-

(continued on page 17)

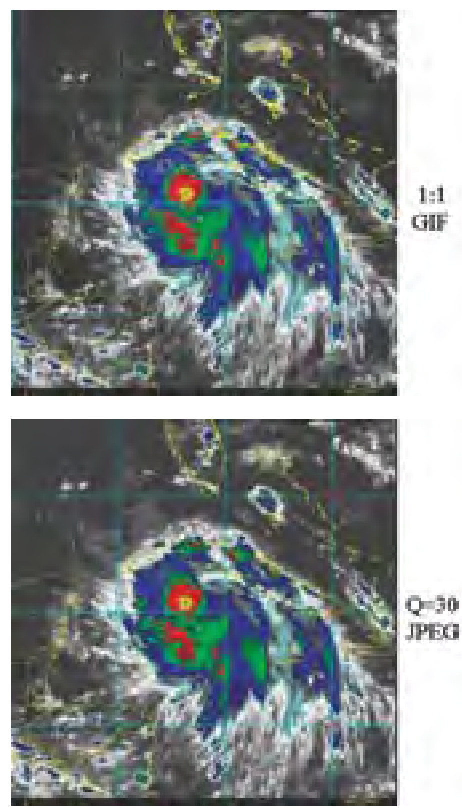

Figure 4, examples of the same image as in Figure 3, but as a GIF image with a size of 88kb (top) and a JPEG image with a quality flag of 30 and a size of 36kb (bottom).

Transmission of Current GOES Imagery *(continued from page 16)*

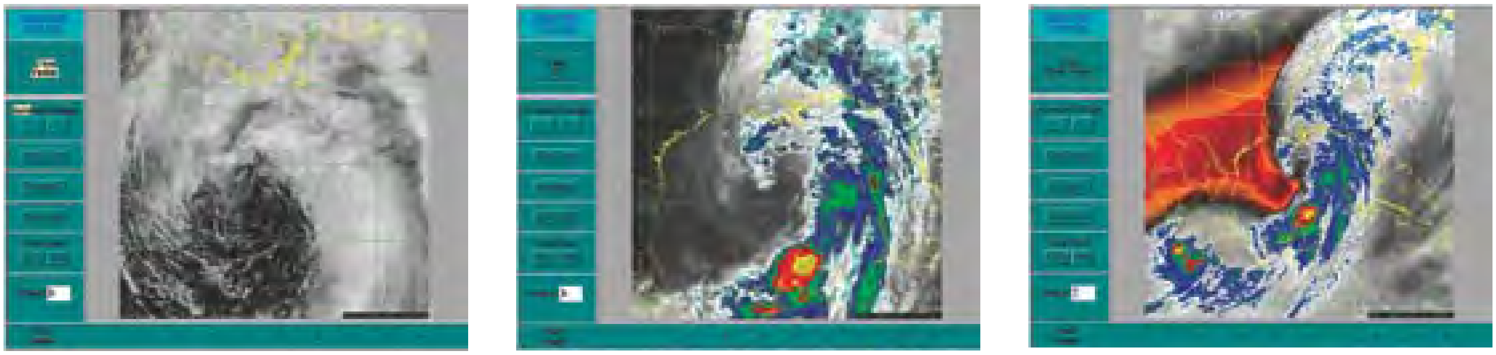

Figure 5, examples of the satellite image products available during Hurricane Isidore on 25 September 25, 2002. The products are 2 km Visible (left), 4km IR (middle) and 8 km Water vapor (right). Note that there are eight images in each product loop. Also this figure displays three different times at 16:45 UTC (visible), 17:15 UTC (IR) and 17:45 UTC (water vapor), respectively.

hourly imagery is displayed for each product (eight frames). The JavaScript also enables the exclusion of bad images and control of the looping speed.

Future Plans

During the 2002 Atlantic Hurricane Season, a system was developed to transmit real-time GOES satellite imagery via SATCOM to the NOAA WP-3 aircraft. First tests of transmission of JPEG images proved successful, but time consuming. Luckily advanced image compression technology in the form of wavelet compression allows better quality satellite imagery to be transmitted in a fraction of the time required for JPEG images. This technology has been utilized to speed the transmission of this imagery, which is severely limited by bandwidth (2400 baud) associated with current SATCOM data transmission. Combining existing software along with these wavelet compression techniques, GOES imagery capabilities now exist to transmit, update, display, and animate real-time GOES imagery on the NOAA WP-3 aircraft.

The display of such data will significantly impact reconnaissance and research missions flown by the NOAA WP-3. System testing will continue through the 2003 hurricane season, and it is anticipated that these capabilities will be fully operational on both aircraft by the 2004 hurricane season.

Acknowledgments:

This project was supported by NOAA grant number NA17RJ1228.

References:

Daubechies, I., 1992: Ten lectures on wavelets, CBMS-NSF regional conference series in applied mathematics, SIAM, Philadelphia, PA, 1992.

Molenar, D.A., K.J. Schrab, and J.F.W. Purdom, 2000: RAMSDIS contributions to NOAA satellite data utilization. Bull. Amer. Meteor. Soc., 81:5, 1019-1029.

Prasad, L., and Iyengar S.S., 1997. Wavelet Analysis with Applications to Image Processing, CRC Press.

Wang, N., and S. Madine, 1998: FX-NET: A Java Based Internet Client Interface to the WFO-Advanced Workstation. 14th Int. Conf. on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology, Phoenix, AZ, amer. Meteor. Soc., 427-429

CloudSat Data Processing *(continued from page 11)*

a Web-based data distribution system that is similar to the type you might use to order a book from a Web-based book store.

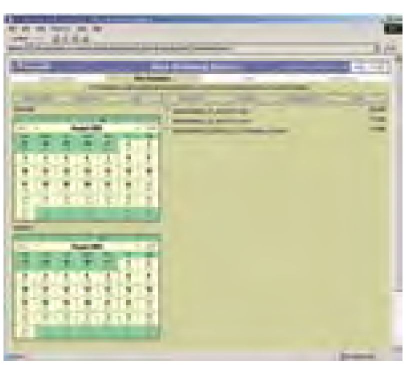

Figure 3, screen shot of the CloudSat Data Distribution System Website. This screen is used to select the date or range of dates that a user is interested in.

The DPC data distribution system will allow the science user to request data by date, time, location, product, and most importantly, volume. Figure 3 shows a sample screen shot of the user interface for this system.

The user selects the product they are interested in, individual data fields from that product, and the date and time range. In addition, they can specify whether they want data for an entire orbit, or – as is most often the case – for a specific location or a radius from that location. Users can store their requests for recall at a later time, and can also set up a schedule to pull their specific data "sector" from every future orbit that passes within the boundaries that they have set. Data will be placed into a directory that is specific to each user and they will be given a specific amount of time to retrieve it before it is removed from the system.

If you would like more information about the implementation of the Data Processing Center and the technical issues discussed in this short overview, please contact Don Reinke at reinke@cira.colostate.edu.

Additional information about the CloudSat mission may be found at http://cloudsat.atmos.colostate.edu

References

Graeme L. Stephens, et al., 2002: The CloudSat Mission and the A-Train. Bull. Amer. Meteor. Soc., 83, 1771-1790.

Grid Computing for Improved Technology Transition

By Andrew S. Jones

Grid computing solves large computational problems in a collaborative fashion by splitting a large problem into smaller parts for processing by each individual computer in the grid. The collaboration can be several fold: such as computing grids that divide up the CPU intensive jobs into smaller jobs across the grid, data grids that divide up storage requests and access across the grid to minimize data bottlenecks and single points of failure, and collaboration grids that tie individual users into a virtual common computing environment to allow collaboration on a common problem, appearing to the end user as a large single system with many on-line users sharing in the collaborative work. Grid technology is currently in a rapid stage of development, with several standards in progress to bring the technology to maturity. At this time, most grid computing systems are highly customized, but are also very powerful tools to perform work that makes the capabilities of the large computing grid accessible in a manner not previously available.

CIRA is leveraging grid-like computing technologies to address a long-standing NOAA technology transition problem. Previously, moving research into operations required several steps (Figure 1), starting with the initial research concept, a reengineering phase that makes the research operationally feasible, and then an operational implementation phase, which includes a variety of quality control, maintainability, and timeliness issues for the many and diverse data and product customers of NOAA. This process tends to create a highly customized operational environment that is very difficult to migrate or port to other sites, or to out-source component development to off-site vendors. Thus, the research integration engineering challenges are difficult problems to solve.

CIRA has developed a grid computing system called the Data Processing Error Analysis System (DPEAS) (Jones and Vonder Haar, 2002) that encapsulates an operational data processing environment using grid computing techniques. This means that the processing system can scale from a laptop to more than a thousand computers collaborating in a grid. This system was tested at NOAA/NESDIS/OSDPD in a prototype environment using just seven nodes in a NOAA High Performance Computing and Communications (HPCC) project called "Harnessing the Spare Computing Power of Desktop PCs for Improved Satellite Data Processing and Technology Transition," (PI: Ingrid Guch, NOAA/NESDIS/OSD). The results of that project demonstrated success with moving a system from research (at Fort Collins, CO) into a demonstrational operational system (at Suitland, MD) and back again (Figure 2). In particular, data sets and algorithms could be modified and updated with relatively modest levels of effort for the research and operational staff collaborators. An initial cost-benefit analysis from this project based on 100 office PCs versus a dedicated computing cluster, results in cost savings exceeding $1,000,000 over a 5-year period based on a Total Cost of Ownership model. Future studies on technology transition cost savings are planned.

This work demonstrates the growing opportunities of grid computing technologies to significantly streamline the traditional technology transition process, with scientists and operational staff collaborating on a common grid system generating products on under utilized desktop computing resources.

References:

Jones, A. S., and T. H. Vonder Haar, 2002: A dynamic parallel data-computing environment for cross-sensor satellite data merger and scientific analysis. J. Atmos. and Oceanic Technol., 19, 1307-1317.

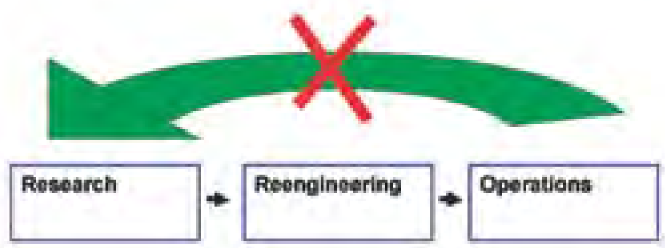

Figure 1: Current satellite technology transition paths can require substantial reengineering. Technology transfer benefits to the research community are minimal.

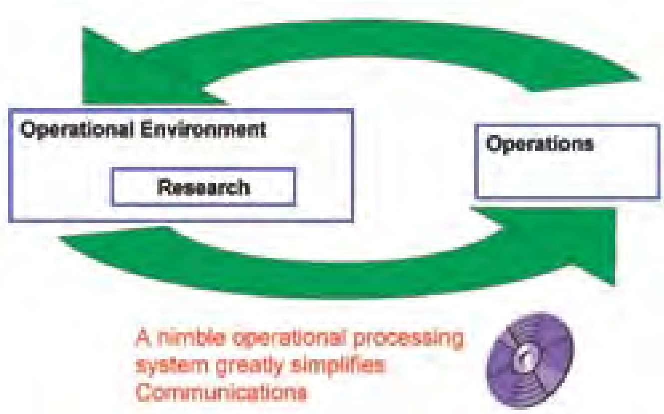

Figure 2: New grid-enabled technology transition paths are now two-way paths and are much simplified. This allows operational benefits to flow back to the research community, and facilitates efforts to replicate the operational system (useful for crisis management planning). The close collaborations result in benefits to both communities.

A Prototype NOAA Computational Grid

By Daniel S. Schaffer

CIRA and the NOAA Forecast Systems Laboratory (FSL) are leading an effort to develop a prototype NOAA computational grid. We have developed a rudimentary grid consisting of a few Intel Linux machines located at the Pacific Marine Environmental Laboratory (PMEL) and several others situated at FSL. The nodes are linked using the Globus toolkit, a middleware package that enables the nodes to be treated as if they all belong to the same machine. Communication between the PMEL and FSL nodes occurs over the Abilene Next Generation Internet (NGI) which links FSL and PMEL.

In order to analyze the behavior and gauge the performance of the grid, a version of the Weather Research and Forecast (WRF) model has been coupled with the Regional Ocean Modeling System (ROMS). In the experiment, the ROMS model is run on the PMEL nodes and the WRF model runs on the FSL nodes. Boundary conditions such as wind stress and sea surface temperature are exchanged between the models periodically (Figure 1).

Since the WRF and ROMS models execute on different grids and each model is running on a parallel computer, the issue of how to exchange boundary conditions becomes important. In short, a parallel re-gridding operation is required. A Department of Defense funded project has provided a software library to facilitate this re-gridding. The library consists of a top-level Application Program Interface (API) in the form of the WRF Input/Output API, and a back end constructed with the Model Coupling Toolkit (MCT) provided by the Argonne National Laboratory. The WRF/ROMS boundary condition exchange can then be carried out with this package.

A test problem was executed that simulates the evolution of a super-cell off the coast of Florida during summer (Figure 2). ROMS ran with 482x482x15 grid points and a time-step of 240 seconds. WRF ran with 150x150x20 grid points and a time-step of 40 seconds. To measure the performance of the boundary condition exchange, an integration consisting of 15 WRF time-steps, and 2 ROMS time-steps was executed. The WRF model ran on two CPUs of a single node at FSL, while the ROMS model ran on both CPUs of a PMEL node. Boundary conditions were exchanged three times during the integration. The integration time for WRF was 64 seconds – 1.3 of which was spent communicating over the grid. The ROMS integration required 57 seconds – 0.8 of which was cross-grid communication. Thus, it is clear that coupling over the grid in this fashion is feasible since even if each model were run on 32 processors (and assuming 75% parallel efficiency), then the model run-times (2.7 and 2.4, respectively) would be longer than the cross-grid communication time. Moreover, if the coupling were to occur over the TeraGrid, which has much higher bandwidths, then the observed communication times would be reduced significantly.

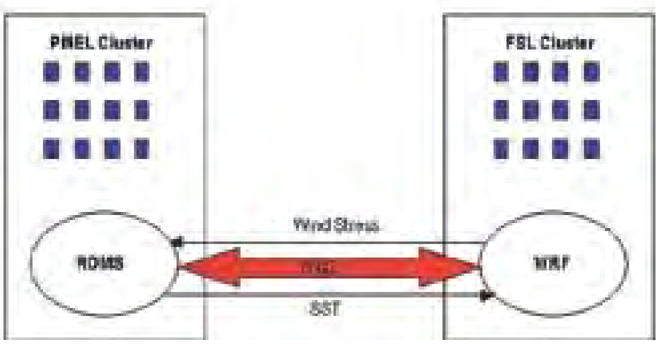

Figure 1, parallel coupled model running on a prototype NOAA grid. The Weather Research and Forecast (WRF) model runs on a cluster of Intel Linux nodes located at the NOAA Forecast Systems Laboratory (FSL). The Regional Ocean Modeling System (ROMS) runs on a Linux cluster at the Pacific Marine Environmental Laboratory (PMEL). The exchange of boundary conditions occurs over the next generation Internet (NGI).

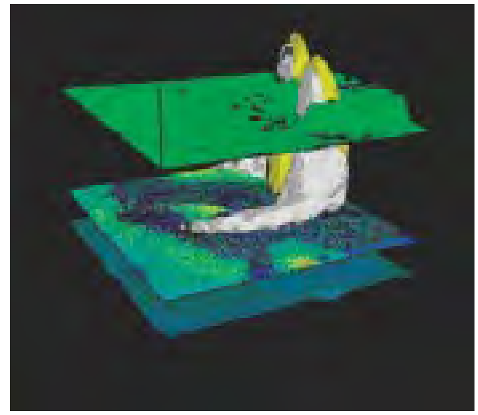

Figure 2, super-cell simulation off the coast of Florida as simulated by integration of a coupled WRF/ROMS model running over a grid.

The importance of this result can be seen in light of the following hypothetical example. Imagine that a user needs 100 CPUs to run a coupled model in a desired period of time. Suppose 50 are accessible at each of two geographically separated clusters. Then only by running one model on one cluster, the other model on the second cluster and coupling over the grid would it be possible to complete the integration on time.

The ultimate objective of this research is, of course, to construct a NOAA computational grid. Ultimately, the physical location of NOAA clusters and individual nodes should be mostly transparent. A NOAA grid user should be able to simply make a request for a certain number of nodes with various physical characteristics. This request would then be routed to any available nodes on the grid meeting those parameters. A proposal for follow-on work has been submitted to NOAA to develop a prototype grid consisting of clusters and individual nodes at FSL, PMEL, and the Geophysical Fluid Dynamics Laboratory (GFDL). The envisioned prototype will include state-of-the-art security mechanisms to allow access by authorized users only. If successful, the prototype will pave the way toward seamless shared access of NOAA computational and data resources.

Data Assimilation *(continued from page 14)*

Foundation, National Aeronautic and Space Agency, and National Oceanographic and Atmospheric Administration.

Current data assimilation projects include:

- 4D variational (4DVAR) data assimilation methodology with CSU Regional Atmospheric Modeling System (RAMS)
- Assimilation of cloudy radiances using visible and infrared satellite measurements
- Radiative transfer modeling for all weather radiance data assimilation
- Microwave land surface modeling for soil moisture data assimilation
- Ensemble data assimilation methodology
- Assimilation of precipitation measurements
- Estimation of model errors in data assimilation

The data assimilation research results produced through these projects by a team of researchers are novel and very encouraging. For example, a Regional Atmospheric Modeling and Data Assimilation System (RAMDAS) was developed and successfully applied in experiments with the conventional meteorological observations and the GOES visible and infrared imager radiances for several mesoscale weather cases. The unique feature of RAMDAS is that it assimilates the conventional observations by the way of the WRF 3DVAR operator originally developed at NCEP. The new data assimilation results show potential for improving high resolution forecasts with these typically synoptic-scale observations but correlated properly with the mesoscales by the 4D dynamical data assimilation method in RAMDAS.

The data assimilation study with the GOES imager data in RAMDAS demonstrates for the first time that cloudy radiances can be readily assimilated into mesoscale weather forecasts and they have the strong potential to provide appropriate initialization of clouds and to improve their prediction. This result points toward improved severe weather forecast skills. The RAMDAS studies also demonstrate the importance of including the estimation of the model error in the data assimilation both as means of relaxing the strong influence of the model when processing the observations and as valuable diagnostics to help improve the forecast model.

Other recent studies not including RAMDAS produced very important results as well. For example, the study on information content of the new satellite microwave soil moisture measurements demonstrates methodology for diagnosing conditions for weak and strong influence of these measurements in the data assimilation.

Research on the use of atmospheric radiative transfer models in data assimilation is another integral part our research at CIRA. Radiative transfer models are needed to link the mesoscale weather forecast with the satellite radiance measurements in the assimilation algorithm. A combined radiative transfer model applicable to the spectrum of visible and infrared measurements has been developed based on the state-of-the-art theoretical modeling studies of the radiative transfer in both cloudy and clear atmospheric conditions.

The challenge extends to the investigation of new assimilation methods. The ensemble data assimilation research in CIRA addresses the problem of the optimal estimation of the complex nonlinear mesoscale weather as well as the evaluation of associated uncertainties. In this way the data assimilation procedure could be made more effective in reducing the most significant errors in the forecast typically caused by nonlinear interactions between different scales.

In summary CIRA's data assimilation research is focused on improving the analysis and prediction of extreme weather by way of incorporating satellite remote sensing and other observations with advanced data assimilation methodology.

Communiqué *(continued from page 13)*

In the Boulder offices, **Brent Shaw** was promoted to Research Associate III for his close collaboration with the NOAA Forecast Systems Lab on various innovative mesoscale modeling and data assimilation research. **Ali Zimmerman** was promoted to Research Associate II for her noteworthy contributions to the highly successful GLOBE Program now managed by the new UCAR/CSU partnership.

Four promotions occurred among the members of the CIRA/RAMM Team. **Bernie Connell** made valuable contributions to international training activities in Central and South America, and is continuing her research on applications of GOES satellite data to cloud formation, and fire and volcanic ash detection. **Louie Grasso** has recently expanded his research to include the development of new products for severe weather and tropical cyclone analysis for future satellite sensors by combining numerical model simulations and radiative transfer models. **John Knaff** continues to be actively involved in the improvement of satellite techniques for operational tropical cyclone analysis and forecasting at the Joint Typhoon Warning Center in Guam, and the National Hurricane Center in Miami. All three were all promoted to Research Scientist II. **Dave Watson** is heavily involved in the effort to move from the UNIX to the LINUX computing environment and the upgrade of Web-based methods for dissemination of experimental satellite products. He was promoted to Research Associate II.

In the Fort Collins office **Michael Hiatt** was promoted to Research Associate IV in recognition of his important work managing the CIRA infrastructure. Michael has developed highly innovative data archival applications and has created a cost-effective computer infrastructure for CIRA. **Adam Kankiewicz** earned a promotion to Research Associate II for his leadership in the Center for Geosciences field programs. **Mary McInnis-Efaw** was promoted to Manager as she took on department manager responsibilities, and **Lance Noble** was promoted to Associate Manager as he took on more responsibility in the management of CIRA finances. **Tomi Vukicevic** was promoted to Research Scientist III for her work with the Center for Geosciences, and her leadership in CIRA's cutting-edge four-dimensional data assimilation research.

Volcanic Ash Coordination Tool Project

By Jim Frimel

The 2001 eruption of Mt. Cleveland in the Aleutian Islands of Alaska provided instances of inconsistent weather advisory products that were generated for adjacent Flight Information Regions. In response, collaborative approaches for generating Volcanic Ash Advisories are being jointly performed and evaluated by the Anchorage Volcanic Ash Advisory Center (VAAC), Alaska Volcano Observatory (AVO), and Anchorage Air Route Traffic Control Center (ARTCC) Center Weather Service Unit (CWSU). The Volcanic Ash Coordination Tool (VACT) is being used at each of these operational units to simultaneously view identical displays and collaborate on weather information in near real-time to help create a suite of fully consistent advisories and forecasts for volcanic ash (Figure 1).

CIRA researchers at NOAA's Forecast Systems Lab (FSL) will extend and investigate the current capabilities of the FX-Collaborate (FXC) and AWIPS systems to include volcanic ash data sets, dispersion models, and tools for generating volcanic ash products so that the goal of creating a consistent set of advisories from different organizations can be achieved. A rule-based approach to establishing guidelines for collaboration will be developed and tested during this project.

The focus of the first release of the VACT, version 0.0, was to train the participating agencies (VAAC, AVO, and CWSU) on the capabilities of the VACT so that feedback from the users could help drive the research efforts in creating a more useful tool. The FXC and AWIPS components of the VACT were enhanced to include satellite volcanic ash dispersion model displays, volcano locations, jet routes, and flight information areas for Kamchatka and the Aleutian Islands. See Figures 2 and 3 for sample satellite imagery displays for the Aleutians and the Kamchatka peninsula.

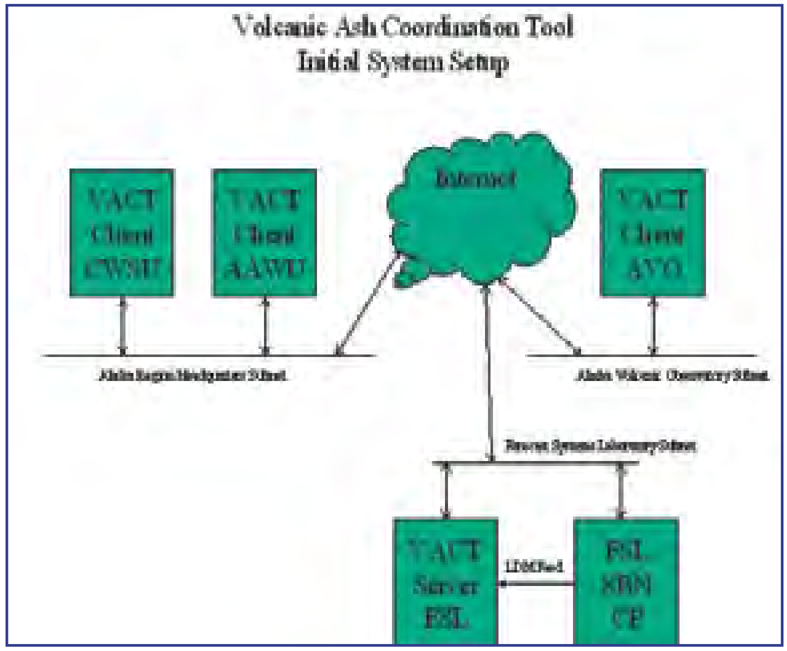

Figure 1, Volcanic Ash Coordination Tool Initial System Setup

Future work will focus on adding additional satellite, radar, and dispersion model data sets to the VACT. A rule-based approach to collaboration will be realized and tools for the forecasters will be developed for outputting their advisories to the public.

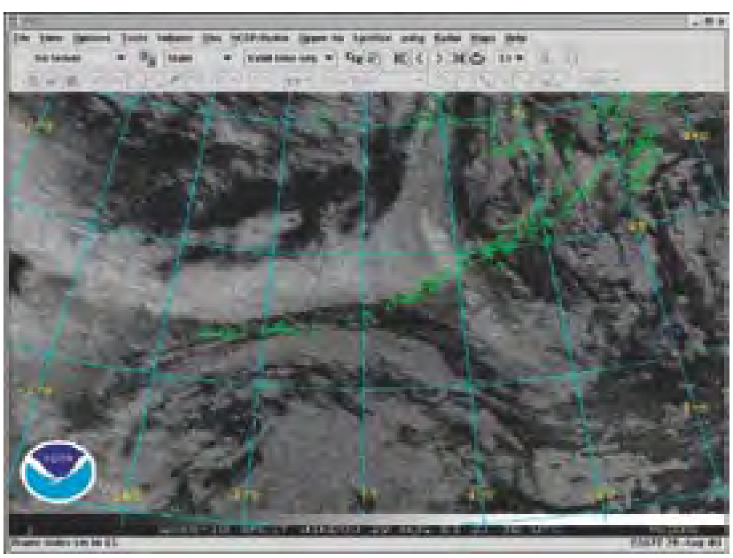

Figure 2, FXC Display of the GOES 10 Satellite Volcanic Ash Dispersion Model Split Window for Aleutian Islands

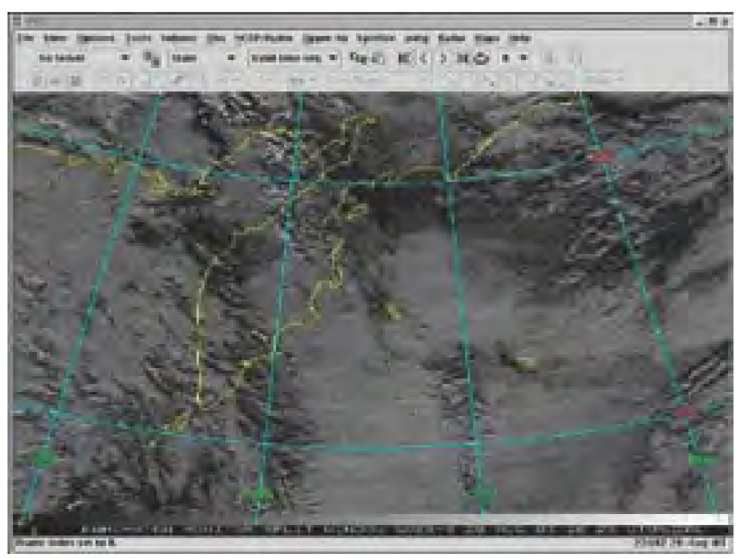

Figure 3, FXC Display of the GOES 9 Satellite Volcanic Ash Dispersion Model Split Window for Kamchatka

The Object Data System

By Paul Hamer

This paper describes the evolution of the NOAA Forecast Systems Laboratory (FSL) Central Facility data ingest, routing and product generation systems from the legacy configuration to a system that is both highly configurable and one that supports the processing of datasets in real- and non-real-time (case-study) situations. In addition, we will highlight the significant cost savings made through the use of Object Oriented Analysis and Design (OOAD) coupled with the use of Open Source[1] for both the development and deployment of the resulting systems. In particular, we will discuss the reduced software development costs in generating the required products from both existing and new datasets obtained by the new system. Finally, we will give examples of how the system supports the FSL mission of technology transfer to entities outside of the laboratory and possible future developments within the Central Facility.

1. Introduction

At FSL, the Central Facility plays an important role in supporting both scientists and engineers in the development of the forecast systems to be used, now and in the future, by both the National Weather Service and private industry. Within the Central Facility, the Data Systems Group (DSG), comprised of CIRA researchers, is responsible for the obtaining, storing, transforming, and distributing of any dataset required by the research groups within FSL.

Since dataset volumes and types are ever increasing, a refactoring of existing ingest and processing was considered necessary in order to reduce both software costs and the software development lifecycle duration. Initial efforts to reduce costs involved the introduction of Open Source software for both development and deployment (e.g. the GNU Compiler Collection (GCC) and other GNU tools.) A switch to Linux running on less costly hardware was also embarked upon while continuing to maintain software portability to other platforms. One of the largest cost reductions can be made by reducing the software development time and therefore refactoring of the legacy systems using OOAD was considered important.

First, a description of the legacy system, developed in the early nineties and called the Posix Data System (PDS) that contained, at its core, the Networked Information and Management Client-Based User Service (NIMBUS)[2]. PDS/NIMBUS was the DSG switch from proprietary operating systems to an open architecture under UNIX. The basic concept was for processes to pass data and notifications via a routing process (called the "cloud") to other processes that had registered to receive those data (see Figure 1).

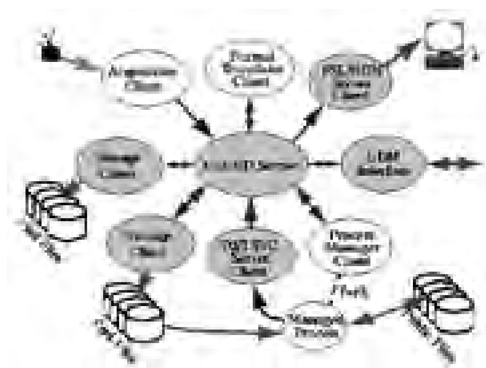

Figure 1, PDS/NIMBUS Architecture

PDS/NIMBUS worked well initially, but the effort required to introduce new datasets became increasingly costly in both software development and the configuration required for the software to operate. The development of the Object Data System (ODS) was initiated with the aim of reducing the time taken to introduce new datasets and generate user requested products while, at the same time, increasing the flexibility of the configuration to allow the rapid introduction of the new software. As an additional benefit, the increased flexibility and open source development tools used enable the software to more easily port and run outside of the confines of the Central Facility and thus support the FSL mission of technology transfer.

The Object Data System Design

From the start, the ODS intended to leverage existing software packages, the Local Data Manager (LDM)[3] from Unidata for example, and to decouple the metadata from the datasets themselves. Further, the design of ODS was primarily driven by looking for known "patterns" of software design for solutions to the problems faced by DSG in carrying out its function of supplying required data.

The first design decision made was to remove the "cloud server" from the distribution scheme. Consequently, data ingest and routing is now to be handled through LDM. The advantages of this approach are that LDM is widely used, several existing software packages are available for handling common meteorological datasets and, being open source, developers can extend LDM to even better support ODS development. The extensions to, and bug fixes in, LDM are discussed with Unidata regularly and, when implemented, introduced into new versions of LDM that are then released to the user community. Since a full description of ODS is beyond the scope of this article, the following sections are intended to highlight some of the concepts used in the development of the system and give an example of how the Central Facility contributes to FSL's mission.

Design Patterns[4]

Most well designed, object-oriented architectures contain patterns, where a pattern is nothing more than a simple and elegant solution to a specific problem in object-oriented software.

Figure 2, the Observer Pattern

Within ODS, we have attempted to identify those patterns and have implemented solutions accordingly. For example, access to data from LDM is achieved through use of the Observer pattern which allows clients to access LDM data in real-time as they arrive in LDM product queue. (See Figure 2.) From this pattern we were able to develop processes to handle all data sets in a common

(continued on page 23)

way and therefore simplify both development and deployment. Other patterns commonly used are the Singleton/Multiton, Proxy and Factory patterns, all of which provided well-crafted solutions for the problems faced in the development of ODS.

ODS Architecture

Using LDM as the data routing mechanism allows ODS to use a facility that provides for the triggering of jobs on events using notifications. This is handled through a product queue action where specified product keys will result in jobs to handle a given event being spawned. (See Figure 3.)

Processes responsible for product generation, typically spawned on arrival of the original data, arbitrate between objects that are proxies for real data types, GRIdded Binary (GRIB), for example, and those objects responsible for certain products, like netCDF. This much simplified architecture supports rapid introduction of new datasets and a more manageable distribution model for the processing of all data.

As an example of the significant savings available under the new system, consider the handling of new model data from a generating center previously unknown to FSL using a generating center-defined parameter table. In terms of developer effort, ODS needs less than a day of person effort to take a requirement to make this new model in GRIB available as NetCDF compared with 15 or more days of effort under the legacy systems to support the same requirement. This is possible because no software is required to be modified or written, tested and integrated; the ODS solution needs only the metadata in the form of the parameter table from the originating center and netCDF description file for the required product. Similar savings have been realized for other data types using the ODS model.

Another significant advantage of ODS is the archive and recall capability. Under ODS, data inserted into LDM is tagged with a product key that includes original receipt time. This has enabled the development of the FSL Data Repository (FDR) for data archive using the same class libraries that ODS uses for real-time processing from LDM data ingest. FDR archives can then themselves be used directly in any ODS configured system to generate case-study data products by replaying data from the archive through a LDM.

Technology Transfer

A core element of FSL's mission is to transfer technology developed within the research environment to entities outside FSL. In support of that mission, an effort has been made to package elements of ODS using GNU's autoconfigure[5] tool. Packages available from DSG currently include the following ODS elements:

LdmNexrad2NetCDF – Level II Nexrad data to NetCDF;

Grib2NetCDF – GRIB editions 1 and 2 to NetCDF;

GOES GVAR Ingest and NetCDF production.

All of these packages are in use outside of FSL.

To more fully illustrate this aspect of FSL's mission, we can cite the distribution of the GVAR ingest and associated NetCDF generation software as an example. Internally, GVAR processing within the Central Facility is handled for all operational GOES satellites in order to support, amongst other things, model initialization and analysis for the Local Area Prediction System (LAPS). The Central Weather Bureau (CWB) of Taiwan use LAPS for local forecasting and, due to the failure of Japan's GMS weather satellite with the subsequent retasking of the GOES-9 satellite to the region while waiting for a Japanese replacement to become operational, needed to deploy a low cost solution for ingesting GVAR data from GOES-9. DSG's solution was chosen because of the low-cost compatibility with LAPS and the ability to deploy the solution quickly into an operational environment.

Members of DSG were able to contribute in the specification of the hardware, deploy, install, and configure ODS software within the required timeframe to support operational start date for GOES-9 in large part due to the ODS architecture and development. We've also used the transfer opportunity to develop new tools to help FSL and the Central Facility to better manage GVAR data internally as these tools were additional requests made by CWB to help address operational concerns. Currently the ODS GVAR is in operation helping CWB provide weather forecasts and information for their general public.

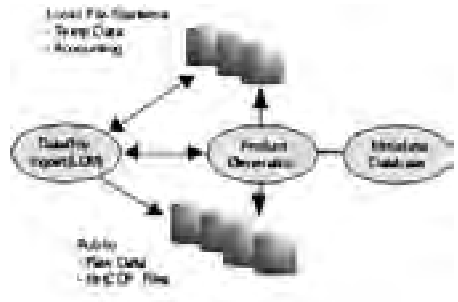

Figure 3 – ODS Architecture

The Future

The Data Systems Group continues to look at new technologies to better serve the laboratory, NOAA, and other groups both within and outside the government.

Perhaps the largest challenge facing the Central Facility, and indeed all of NOAA, is the ever increasing data volume. It is estimated that in 2004, NOAA agencies will generate more data in one year than was contained in the entire archive in 1998. With this growth, the challenge of obtaining and analyzing these data has also grown. DSG is starting work on cataloging real-time and archived data using the latest technologies to both describe and distribute these descriptions as extensible catalogs using the eXtensible Markup Language (XML). The work done under ODS will help support this effort by using the framework in place for the handling of metadata within the architecture described above. This, in turn, will enable DSG to meet future data requirements placed on the Central Facility by FSL and other agencies.

References

[1] Open Source Initiative (OSI) – http://www.opesource.org

[2] Networked Information Management client-Based User Service – http://www-fd.fsl.noaa.gov/nimbus/

[3] Local Data Manager (LDM) – http://www.unidata.ucar.edu/content/packages/ldm/index.html

[4] Gamma, Helms, Johnson and Vlissides - Design Patterns – Elements of Reusable Object-Oriented Software. Reading, Massachusetts: Addison-Wesley, 1995.

[5] Creating Automatic Configuration Scripts – http://www.gnu.org/directory/GNU/autoconf.html

CIRA Mission

The Mission of the Institute is to conduct research in the atmospheric sciences of mutual benefit to NOAA, the University, the State and the Nation. The Institute strives to provide a center for cooperation in specified research program areas by scientists, staff and students, and to enhance the training of atmospheric scientists. Special effort is directed toward the transition of research results into practical applications in the weather and climate areas. In addition, multidisciplinary research programs are emphasized, and all university and NOAA organizational elements are invited to participate in CIRA's atmospheric research programs.

The Institute's research is concentrated in several theme areas that include global and regional climate, local and mesoscale weather forecasting and evaluation, applied cloud physics, applications of satellite observations, air quality and visibility, and societal and economic impacts, along with cross-cutting research areas of numerical modeling and education, training and outreach. In addition to CIRA's relationship with NOAA, the National Park Service also has an ongoing cooperation in air quality and visibility research that involves scientists from numerous disciplines, and the Center for Geosciences/Atmospheric Research based at CIRA is a long-term program sponsored by the Department of Defense.

Cooperative Institute for Research
in the Atmosphere
College of Engineering-Foothills Campus
Colorado State University Fort Collins, CO 80523-1375
(970) 491-8448
www.cira.colostate.edu

If you know of someone who would also like to receive the CIRA Newsletter, or if there are corrections to your address, please notify us.