REINAS: Real-Time Environmental Information Network and Analysis System -- Annual Report 1995

University Research Initiative -- Office of Naval Research, No. N-00014-92-J-1807

Principal Investigator: Patrick E. Mantey, Ph.D.
Jack Baskin Professor of Computer Engineering
Address: Baskin Center for Computer Engineering and Information Sciences
University of California, Santa Cruz
Santa Cruz, CA 95064 U.S.A.
email: mantey@cse.ucsc.edu
telephone: 408-459-2158
fax: 408-459-4829

Co-Principal Investigators:

Wendell Nuss, Ph.D.
Address: Department of Meteorology
Naval Postgraduate School
Code MR/NU
Monterey, California 93943-5000
email: nuss@arashi.met.nps.navy.mil
telephone: 408-656-2308
fax: 408-656-3061

Bruce R. Gritton
Address: Monterey Bay Aquarium Research Institute
Pacific Grove, California
email: grbr@mbari.org
telephone: 408-647-3733
fax:

Darrell D.E. Long, Ph.D.
Address: [Same as PI]
email: darrell@cse.ucsc.edu
telephone 408-459-2616
fax: 408-459-4829

Alex T. Pang, Ph.D.
Address: [Same as PI]
email: pang@cse.ucsc.edu
telephone 408-459-2712
fax: 408-459-4829

Jose Joaquin Garcia-Luna, Ph.D.
Address: [Same as PI]
email: jj@cse.ucsc.edu
telephone 408-459-4153
fax: 408-459-4829

Glen G. Langdon, Ph.D.
Address: [Same as PI]
email: langdon@cse.ucsc.edu
telephone 408-459-2212
fax: 408-459-4829

TITLE OF GRANT: Real Time Environmental Data Management

LONG TERM GOALS: The major goal of this research is to achieve faster and more accurate weather prediction. The ocean and atmospheric processes of both operational and research interest at a regional scale, occur on time and space scales that require nearly continuous measurement, as opposed to discrete time interval measurements as used in the past. In addition, measurements and calculations must be made of both mean and fluctuating values in order to understand both mean transports and properties and turbulent fluxes. Observations are often sparse and irregular and must be assimilated into a consistent representation of the phenomena being observed. Techniques are needed to properly collect these data, assimilate them into analysis and/or forecast models, and display both the data and model output using creative visualization software. The quantity of potential data and diversity of possible models require development of sophisticated processing methods for researchers to intelligently sort out the different scales of processes that are active in the environment.

The focus of this effort is to develop the capability for real time collection, assimilation, and display of in-situ and/or remotely sensed environmental data. The research includes methodologies for data collection, compaction and storage as well as appropriate physical analysis and modeling. Integration of database information systems and visualization with direct manipulation interfaces, and development of novel means for visualizing and melding large data sets comprised of diverse data types are major goals.

SCIENTIFIC AND TECHNICAL OBJECTIVES

From a scientific perspective, the primary objective of this research is to develop a system to support nowcasting at a regional scale by resolving smaller space/time scales of the phenomena. This requires understanding of the processes being observed, and the corresponding requirements for a system to support their real-time monitoring. The Real-Time Environmental Information Network and Analysis System (REINAS) is an engineering research project with the goal of designing, developing and testing an operational prototype system for data acquisition, data management, and visualization to support the real-time utilization of advanced instrumentation in environmental science. .

Specific technical objectives for the REINAS system are that it:

- Support interactive user access to real-time and retrospective data,
- Provide the performance required to support the most frequently requested products and services,
- Provide access to resources and devices through a common data model supporting rapid system configuration,
- Support dynamic control of devices for real-time interactive scientific investigation,
- Provide fault tolerant data collection, avoiding data loss due to communication link failures, and
- Provide security features that support restricting user access and control with respect to data and equipment.

APPROACH

The project is a multi-year effort of the Baskin Center for Computer Engineering and Information Sciences of the University of California, Santa Cruz (UCSC), in cooperation with environmental scientists from the Naval Postgraduate School (NPS), and Monterey Bay Aquarium Research Institute (MBARI). The REINAS system has been designed for regional real-time environmental monitoring and analysis. REINAS is designed to provide on the desk top:

- A set of tools to configure and collect data from instruments in the field in real time
- An integrated problem solving and visualization system supporting individual and collaborative research using both historical and modeled data, and
- A logically consistent distributed data base that stores data in the data base independent of file format and which maintains metadata describing where and how data was obtained (the data base tracks data pedigree).

Professionals in the environmental sciences will have the ability to observe, monitor, and analyze regional oceanographic and meteorological phenomena from their desk top. REINAS serves different user groups: ***Operational forecasters*** monitor current conditions, view standard data products, synthesize new data views, and issue forecasts and warnings. ***Modelers*** analyze new model products and compare them with other models and with past and present conditions. ***Experimental scientists*** collaborate with other scientists on-line, observe individual data fields as they are collected, and may modify data collection methods for an on-going experiment. Finally, ***instrument engineers*** add new equipment to the system, access metadata describing individual devices and methods of calibration, study maintenance records, and profile sensor quality.

The systems software design for REINAS involves development efforts in several areas. Meteorological and oceanographic instruments attached to the REINAS system are augmented with microcomputers to become “intelligent instruments” and attached as Internet nodes to the system. A network of such instruments has been developed in the Monterey Bay region, employing some new instruments and some existing instruments belonging to the participating institutions: Naval Postgraduate School and MBARI. A set of load paths is used to move data from collection devices into the system. A commercial (relational) database system is used as a central component of REINAS, augmented with software implementing a memory-resident cache to hold temporarily the voluminous data from real-time instruments, and to insure access of current data by real-time applications and users. A set of control applications is used to deploy, configure, and steer collection devices, and, set of data access tools supports basic system administration. Access to REINAS functionality is provided by a set of Application Programmer Interfaces (APIs). Visualization software supports the integrated display of data via a rich set of visualization tools which couple directly (via an API) to the database.

REINAS contributes to scientific productivity by eliminating the need for each researcher to invent their own situation-specific access routines. The REINAS load paths effectively implement automated, real-time methods to load data into the database, as opposed to the manual loading methods often used in conventional scientific databases. REINAS load paths also perform any data parsing that is required during the first steps of the data load process. The use of a relational database with a special API in REINAS provides a common data access model for several different database products. Thus, researchers who implement a scientific application using the REINAS API effectively have direct access to the data with which they work in common binary format. Unlike file specific applications, REINAS applications should continue to be of value regardless of the specific data feed, database, or machine operating system for which the necessary REINAS software exists.

TASKS COMPLETED

A significant milestone was achieved in the spring of 1995 when <u>**REINAS went on-line with its first real-time source**</u>, the Long Marine Lab wind profiler. Following some testing and verification, REINAS began providing simple support for the work of oceanographic and meteorological users. This support has included the use of regular meteorological data products generated from data in the database, Codar data obtained from the database with specially developed SQL queries, and Santa Cruz wind profiler data acquired from the real-time cache across the Internet. The current REINAS implementation now consists of approximately 165,000 lines of code which run across several hardware platforms.

The REINAS <u>**database design**</u> supports measurement data in a temporal / spatial organization, and also includes support for metadata (i.e. the data about the instruments, calibration, etc.) within the database. During this past year, <u>**load paths**</u> have been created (software) to automatically and continually transfer data into the data base from a variety of instruments. Any instrument whose data comes directly into REINAS is augmented by a microcomputer (an Intel-based PC running Unix) that is Internet accessible at the instrument site. Data collected by <u>**instruments**</u> are fed to the PC serially where REINAS system software formats and logs the data appropriately, and ships the data to the database. Data that does not come directly to the REINAS system may come from a number of different paths (or ``virtual" instruments), such as archival files, UNIDATA, or anonymous FTP. The types of instruments that are currently connected directly to the REINAS system are standard research quality meteorological (MET) stations, CODAR (Coastal Ocean Dynamics Applications Radar) high-frequency (HF) current measuring radars, wind profilers, and CTD's (Conductivity, Temperature, Depth sensors), although the system is not limited to these types of instruments. Any instrument that can provide a serial output stream or file can potentially be connected to the REINAS system in a realtime fashion.

The <u>**standard REINAS MET station**</u> consists of a number of individual sensors that include wind speed, wind direction, air temperature, barometric pressure, and humidity. Additional sensors that some sites are equipped with are rain gauge sensors and solar irradiance sensors. The outputs of these sensors are hooked to a data logger (Campbell Scientific CR-10/CR-21 data loggers are the standard ones used by REINAS). Currently, there exist approximately eleven MET stations that are connected directly to REINAS. Of these, two are deployed on moving vessels. One is currently aboard the vessel R/V Pt. Lobos and is Internet accessible via their existing network shipboard microwave connection. The other has recently been deployed aboard the NOAA research vessel R/V McArthur and was Internet accessible via a UHF radio link from UC Santa Cruz to the MET station aboard the vessel. One of the other MET stations in the realtime network was built on a trailer and is easily deployable to a location of interest.

Wind profilers provide a measurement of the horizontal winds as a function of height to a maximum height of approximately 4 km. A realtime link was established via REINAS between two wind profiler locations in the Monterey Bay area, one at Santa Cruz and the other at Point Sur. A realtime link to a wind profiler results in moments (from which wind vectors may be computed) that are received approximately each minute. Receipt of the data in this fashion results in the capability to estimate the wind speed and direction vs height on a much more regular basis than the hourly data provided by the NOAA wind profilers. Two other wind profilers exist in the Monterey Bay area (both of which are at Fort Ord and are owned by the Navy Postgraduate School). Realtime data from these instruments will become available when a network link is established at Fort Ord.

CODARs are a version of high-frequency radar manufactured by CODAR Ocean Sensors, Ltd. Three CODAR instruments, two of which are a newer version called SeaSonde, are deployed in the Monterey Bay area and provide data via REINAS. The two newer units are deployed at Long Marine Lab (Santa Cruz) and Point Pinos, which is at the southern tip of Monterey Bay. These newer units provide measurements of radial ocean surface currents on an hourly basis. An older CODAR unit exists at the MBARI Moss Landing facility and it provides measurements of ocean surface currents every two hours. Data from two or more of these CODAR sites are combined hourly at UC Santa Cruz, which provides maps of ocean surface currents with a resolution of approximately 3 km. HF radar systems measure radial ocean surface currents by measuring the Doppler shift associated with radar energy that is scattered off of ocean gravity waves that are between 5 and 50 meters in length. The difference between the Doppler shift of the returned signal and the expected Doppler shift of the ocean gravity wave allows an estimate of the value of the advection of the waves by the ocean surface current over the patch of ocean that the radar observes. Data from the CODAR units is copied to the REINAS system hourly, where it is then combined to form the vectors.

Other instruments that are part of the current data feed into REINAS include CTD's, ADCP (Acoustic Doppler Current Profiler), thermistor chains, and assorted biochemical sensing instruments, such as those that measure Chlorophyll-A or trace gases. Instruments that make these measurements exist both along the coast (CTD at Granite Canyon) and (more generally) offshore, such as on the MBARI buoys M1 and M2.

During the past year, an investigation into the development and implementation of **remote video** streams as real-time REINAS instruments was initiated. At a system level, this included defining a prototype video instrument, building or acquiring the necessary hardware and software, deploying the camera at an appropriate location, and commencing work on integrating the resulting video data into REINAS. The prototype steerable video platform was defined during the fourth quarter '94, built during the first quarter of '95, and initially deployed at Long Marine Lab (an existing REINAS instrument site) on March 30, 1995. In August '95, it was moved to a more appropriate site (atop a 10 story beach-front hotel) offering a better view of Monterey Bay. The high data-rates achievable with a remote camera instrument (64 kbps to 8 Mbps depending on frame rate, content, and other factors) also motivated a re-examination of the radio modem technology used to network remote sites. A new option augmenting the existing 9 kbps technology was discovered integrated into the REINAS/PC design; this technology provides IP-level data rates exceeding 0.8 Mbps.

The prototype camera platform is currently supporting ongoing dissertation-level research in digital video transport algorithms tailored specifically to meet the unique challenge of a wireless environment and the remote-sensing instrument application. Popular and existing digital video protocols (MPEG I and II, H.261, etc) were designed to transport high frame rate entertainment and teleconferencing oriented imagery over fixed bandwidth, zero-loss connections. Adapting these protocols to work over lossy, variable bandwidth wireless links, in an application where high frame rates are not necessary is unwise and does not work well. As part of the REINAS project, new video protocols are being developed that are better suited for the wireless environment. In addition to research in wireless digital video protocols, the camera platform also provides other unique instrument-level challenges to the REINAS effort. It is the only real-time instrument also capable of providing high-bandwidth input streams into the database. As the camera platform is remotely maneuverable and various camera functions are remotely controllable, this video platform also serves as a prototype steerable instrument that can be manipulated in response to external regional data.

Network development. Communication today among REINAS components and users is being accomplished via a mixed-media networking infrastructure that encompasses new and existing telephone lines, Internet connections, and (point-to-point) radio links, as well as the networking software and hardware needed to control and manage the interconnection of REINAS sites. REINAS requirements include connection with several mobile sites (e.g., boats and trailers with MET stations) and communications support for applications that require the transport of large amounts of information, specially for remote visualizations, where one minute animation can require in the order of 160 Mbps. The information exchanged in REINAS will include multiple media (text, voice, images, graphics and animation, and even video), and such information has to be distributed in real time (e.g., during a multimedia conference among multiple sites) over different types of transmission media, including radio links and high-speed lines. Furthermore, the networking infrastructure of REINAS should allow a very large number of sensors to be incorporated into the system. Accordingly, we have identified six major long-term networking requirements in REINAS:

- The ability to transport multimedia data in real time.
- Scalability to a large number of geographically-dispersed sensors.
- Mobility of sites.
- Fault tolerance.
- Efficient use of multiple transmission media.
- Connectivity to the Internet.

The marked differences between REINAS networking needs and traditional networking technology indicated the need for new multimedia networking solutions to REINAS's unique characteristics and to allow REINAS system engineers to better manage and monitor communication resources in support of data management, data visualization, and user communication. The networking requirements of REINAS translate into the following research problems in Mixed-Media Networking:

- Quick recovery after failures
- Scaling to hundreds or thousands of nodes (sensors, workstations, routers)
- Supporting multimedia traffic over mixed transmission media
- Supporting mobility of nodes (boats, aircraft, terrestrial vehicles, hand-held devices)
- Conserving power and bandwidth (mobile and static nodes)
- Supporting multi point communication with reliable and unreliable data delivery
- Managing mixed transmission media intelligently.
- Supporting distributed applications.

Under REINAS support , we have developed **new algorithms and protocols for channel access**, routing, multicasting, and floor control of distributed applications. We have been developing a new type of channel access discipline called Floor Acquisition Multiple Access (FAMA). FAMA consists of both carrier sensing and a collision-avoidance dialogue between a source and the intended receiver of a packet. Control of the channel (the floor) is assigned dynamically to at most one station at any given time, and this station is guaranteed to be able to transmit one or more data packets to different destinations with no collision with transmissions from other stations. As it is well known, CSMA degrades substantially when the sender cannot sense collisions. Our results indicate that variants of FAMA-NTR can avoid this problem and achieve much higher throughput than CSMA with hidden terminals, and comparable or higher throughput otherwise. Higher throughput can be achieved by increasing the packets sent for each successful floor acquisition.

We have developed a **new type of routing algorithms**, called link-vector algorithms (LVA). An LVA diffuses link-state information selectively based on the distributed computation of preferred paths. We have analyzed the performance of LVA by simulation; our results show that, on the average, an LVA is faster and requires less traffic overhead than the distributed Bellman-Ford algorithm (used in RIP, for example) and any link-state algorithm based on flooding (which is the type of algorithm used in OSPF and IS-IS). We have also developed the first known loop-free routing algorithm based on vectors of distances and second-to-last hops to destinations. We call this algorithm the loop-free path-finding algorithm (LPA). We have analyzed the performance of LPA by simulation; our results show that LPA has better performance than any link-state algorithm based on flooding (which is the type of algorithm used in OSPF and IS-IS) and the most efficient loop-free routing algorithm previously known (the Diffusing Update Algorithm). We are currently implementing our channel access and routing protocols on PCS and these PCS will be soon used as ``wireless routers'' in REINAS.

In the **Visualization** component, during the past year we have completed implementation of the major ideas behind Spray

Rendering. Spray is an extensible framework for visualization. Its current implementation is divided into three modes: monitor, forecast, and analysis. Monitor mode allows users to query environmental sensors for realtime and historical data. Forecast mode generates standard and customized products from model output. Analysis mode allows users to performs more sophisticated analysis on the data.

Over the past year, a major visualization effort was focused on the redesign of Spray with the objective of creating a towards a modeless system. The new version is nicknamed "PET slvg" and incorporates the same modularity as in Spray. PET stands for Products, Elements, and Tools. These are the three different hierarchies that users will encounter. At the top level, users generate visualization products. Products are made up of one or more Tools, usually with the input field assigned. For example, one may generate a Product that is composed of a Tool to generate contours of a vorticity field and overlaid with the output from another Tool that displays registered satellite images. Tools themselves are also modular and are made up of one or more Elements. Elements are simpler computational blocks that operate on a portion of the data set or the entire data set. Individual Elements may be simple (e.g. perform normalization of data) or complex (e.g. perform Fourier transform). An Element template is provided to facilitate adding new Elements and to allow easy integration with other Elements to form new Tools.

During the past year, a new effort was undertaken on **Uncertainty Visualization**. Almost all data has some form of uncertainty -- e.g. instrument calibration errors, sampling errors, noise, model simplification, sensitivity to initial and boundary conditions, approximation/interpolation errors. These uncertainty information must be presented together with the data to give users the ability to localize and distinguish regions where data quality is poor, and eventually make better analysis. We are approaching the problem from four fronts: designing uncertainty glyphs for vector fields, displaying geometric uncertainty in different surface interpolants of scattered data sets, displaying errors in 3-D surface attributes arising from radiosity calculations, and development of new volumetric interpolations using iterated function systems (fractals) whose characteristics are set from the input data. This work is being extended to analyze differences generated by different flow visualization methods (e.g. streamlines, ribbons, streaklines, time lines, etc.). to aid data assimilation and ensemble forecasting.

Results in **Data Compression** have impact in several dimensions of REINAS. Integration of compression transparent to the user is a part of the communications software for the instrument network, the database system design, and for Internet access to database for data of interest. This means compression of raw data on the way into, and processed data on way out of the database. Compression is being applied to a range of instrument data: wind vectors, temperature, rainfall, wind profiles, ocean surface current vectors, ocean salinity, pressure, AVHRR and GOES satellite images, and output of video cameras. Data generated by models is also being compressed. For visualization, compression is applied to the raw data, geometric primitives, or rendered screens. A recently initiated activity is considering compression of responses to SQL database queries.

SCIENTIFIC RESULTS

Details of the use of REINAS in meteorology are contained in the report of (Co-PI) Wendell Nuss.

In addition, short deployments of instruments which were attached to the REINAS system and which provided data that could be analyzed in combination with that from the instruments feeding data continuously to REINAS have provided some exciting preliminary results. An example is the 5-day deployment of the dual-site Marconi OSCR HF radar in conjunction with the CODAR instruments. This deployment demonstrated the advantages of having the data from numerous in situ instruments at hand (such as the MBARI ADCP on buoy M1) when comparing data collected by radar instruments. Access to the REINAS database for data collected by the OSCR instruments during the period of the demo revealed correlations of approximately 0.71 between the radial component of the ADCP and the OSCR at Long Marine Lab during the demo period. Furthermore, comparisons between wind direction derived from OSCR radar data at both the Long Marine Lab site and the Moss Landing site to wind direction data collected at the M1 mooring revealed interesting correlations at both sites that appeared to be a strong function of the alignment of the OSCR antenna array relative to the wind direction.

A 175 page REINAS ***Systems Software Manual*** was prepared during the spring and summer of 1995. This manual contains several sections which will be of great use for others wishing to install and use REINAS. Sections of interest include guides for REINAS installation and internals, for REINAS core system APIs, and for writing REINAS device managers.

SIGNIFICANCE: REINAS provides a system architecture and implementation to support real-time operational applications in environmental monitoring and nowcasting. It is easily extensible to new instruments, expandable (by increasing system resources or by creating new instances of the database) and portable (i.e. its design and implementation do not depend on its current applications.) Visualization tools developed for REINAS link the database and scientific visualization, and provide powerful techniques for data assimilation and interpretation.

PUBLICATIONS

P - H.-P. Dommel and J.J. Garcia-Luna-Aceves, ``Floor Control for Multimedia Conferencing and Collaboration," accepted for publication in ACM Multimedia Systems Journal.

P - J.J. Garcia-Luna-Aceves and Shree Murthy, ``A Path Finding Algorithm for Loop-Free Routing," accepted for publication in IEEE/ACM Transactions on Networking.

P - J.J. Garcia-Luna-Aceves and Jochen Behrens, ``Distributed, Scalable Routing Based on Vectors of Link States," IEEE Journal on Selected Areas in Communications, October 1995.

P - Darrell D. E. Long, Bruce R. Montague and Luis-Felipe Cabrera. ``Swift/RAID: A Distributed RAID System," Computing Systems, vol. 7, no. 3 (1994), pp. 333-359.

P - Shree Murthy and J.J. Garcia-Luna-Aceves, ``An Efficient Routing Protocol for Wireless Networks," to appear in Wireless Networks, Special issue on Routing in Mobile Communication Networks, 1996.

P - Alex Pang, ``A Syllabus for Scientific Visualization", in Scientific Visualization in Mathematics and Science Teachings, edited by David A. Thomas, AACE, 1995, pp. 261--283.

P - Alex Pang and Michael Clifton, ``Metaphors for Visualization", Visualization in Scientific Computing '95, R. Scateni, J. van Wijk, P. Zanarini, editors, Springer, pp. 1-9, 149-150.

P- Alex Pang, ``Spray Rendering", IEEE Computer Graphics and Applications, September 1994, pp. 57--63.

P - Craig Wittenbrink, Alex Pang, and Suresh Lodha, ``Glyphs for Visualizing Uncertainty in Environmental Vector Fields", submitted to IEEE Transaction on Visualization and Computer Graphics, 1995.

PA - Alex Pang and Craig Wittenbrink, ``Spray Rendering as a Modular Visualization Environment", Computer Graphics, Vol.29, No.2, FOCUS: Modular Visualization Environments, Past, Present, and Future, edited by Gordon Cameron, 1995, pp. 33--36.

PS - Daniel M. Fernandez, J. F. Vesecky, and C. Teague, "Measurements of Upper-Ocean Current Shear with High-Frequency Radar," Journal of Geophysical Research, accepted pending revision.

PS - Paduan, Jeffrey D., and Rosenfeld, Leslie K., "Remotely Sensed Surface Currents in Monterey Bay from Shore-Based HF Radar (CODAR)," Journal of Geophysical Research, submitted.

CI - Patrick E. Mantey, ``A View of Visualization: Its Origins, Development and New Directions," Keynote Address, Conference on Visualization, SPIE Proceedings on Visual Data Exploration and Analysis, Volume 2178, pp.~2--11, February 1994.

C - Fernandez, D. M., "Remote Sensing Ocean Currents in Monterey Bay with OSCR and CODAR," (formal talk) Eastern Pacific Ocean Conference, Stanford Sierra Camp, Fallen Leaf, CA, Sept. 26-29, 1995.

C - Fernandez, D. M., Vesecky, J. F. , and Teague, C. C., "Measurements of Upper-Ocean Surface Currents with High-Frequency Radar," (formal talk) Proceedings of the IEEE Fifth Working Conference on Current Measurement, St. Petersburg, Fl., February, 1995.

C - Fernandez, D. M., "The REINAS System for Real Time and Retrospective Environmental Measurements," (poster presentation), AGU Fall Meeting, San Francisco, CA, December, 1994.

CP - Darrell Long, Andrew Muir and Richard Golding, ``A Longitudinal Study of Internet Host Reliability," Proceedings of the Symposium on Reliable Distributed Systems, Bad Neuenahr, Germany: IEEE, September 1995, pp. 2-9. (formal talk)

C - Darrell D. E. Long, Patrick E. Mantey, Craig M. Wittenbrink, Theodore R. Haining and Bruce R. Montague ``REINAS: the Real-time Environmental Information Network and Analysis System," Proceedings of the IEEE Computer Society CompCon, San Francisco: IEEE, March 1995, pp. 482-487. (formal talk)

CP - Thomas M. Kroeger and Darrell D. E. Long, ``Predicting Future File-System Actions from Prior Events," Proceedings of the 1996 Usenix Winter Technical Conference, San Diego: Usenix, January 1996, to appear. (formal talk)

CP - Chane L. Fullmer, Darrell D. E. Long and Luis-Felipe Cabrera, ``Performance Guarantees on ATM Networks," Proceedings of the International Conference on Computers and Communication, Phoenix: IEEE, March 1994, pp. 290-296. (formal talk)

CP - Shree Murthy, Cheng Tang and Darrell D. E. Long. ``Performance Guarantees on ATM Networks," Proceedings of the International Conference on Computers and Communication, Phoenix: IEEE, April 1994, pp. 111-114. (formal talk)

C - Alex Pang and Adam Freeman, ``Methods for Comparing 3D Surface Attributes", accepted to SPIE Proceedings on Visual Data Exploration and Analysis, 1996.

C - Craig M. Wittenbrink, G. Fernandez I Ubiergo and G. Langdon, Jr., ``Feature extraction of clouds from GOES Satellite Data for Integrated Model Measurement Visualization", accepted to IS&T/SPIE Symposium on Electronic Imaging: Science and Technology, Image and Video Processing 1996, Robert Stevensons and M. Ibrahim Sezan.

CP - Suresh Lodha, Bob Sheehan, Alex Pang, and Craig Wittenbrink, ``Visual Comparison of Scattered Data Interpolants", submitted to Graphics Interface'96.

CP - Craig Wittenbrink, ``IFS Fractal Interpolation for 2D and 3D Visualization", IEEE Visualization 1995 Conference Proceedings, 1995, pp. 77--84. (formal talk)

C - Alex Pang, ``Visualization Support for Collaborative Spatial Decision-Making", National Center for Geographic Information and Analysis Workshop. Initiative 17: Collaborative Spatial Decision-Making, to appear. (formal talk)

C - Craig M. Wittenbrink, Eric Rosen, Alex Pang, Suresh K. Lodha, and Patrick Mantey, ``Realtime Database Support for Environmental Visualization", Second Workshop on Database Issues for Data Visualization, Visualization '95, Atlanta, GA, 1995. (formal talk)

C - Glen Langdon, Jr., Alex Pang, Craig M. Wittenbrink, Eric Rosen, William Macy, Bruce R. Montague, Carles Pi-Sunyer, Jim Spring, David Kulp, Dean Long, Bryan Mealy, and Patrick Mantey, ``Compression Research on the REINAS Project", NASA, Science Information Management and Data Compression Workshop, Greenbelt, MD, on October 26-27, 1995, pp. 66-73. (formal talk)

C - Alex Pang, Craig M. Wittenbrink and Tom Goodman, ``CSpray: A Collaborative Scientific Visualization Application", SPIE Proceedings on Multimedia Computing and Networking, Vol.2417, 1995, pp. 317--326. (formal talk)

C - Alex Pang and Naim Alper, ``Bump Mapped Vector Fields", SPIE Proceedings on Visual Data Exploration and Analysis, Vol.2410, 1995, pp. 78--86. (formal talk)

C - Craig M. Wittenbrink, Elijah Saxon, Jeff J. Furman, Alex Pang and Sureh Lodha, ``Glyphs for Visualizing Uncertainty in Environmental Vector Fields", SPIE Proceedings on Visual Data Exploration and Analysis, Vol.2410, 1995, pp. 87--100. (formal talk)

C - Alex Pang and Naim Alper, ``Mix & Match: A Construction Kit for Visualization'', IEEE Visualization 1994 Conference Proceedings, 1994, pp. 302--309. (formal talk)

C - Alex Pang, Jeff Furman and Wendell Nuss, ``Data Quality Issues in Visualization'', SPIE Proceedings on Visual Data Exploration and Analysis, Vol.2178, 1994, pp. 12--23. (formal talk)

C - Alex Pang and Dan Fernandez, ``REINAS Instrumentation and Visualization'', Proceedings of Oceans'95, San Diego, California, October 1995, pp. 1892--1899. (formal talk)

CP - H.-P. Dommel and J.J. Garcia-Luna-Aceves, ``Floor Control for Networked Multimedia Applications'', ACM SIGCOMM'95 Workshop on Middleware, Cambridge, MA, August 28-September 1, 1995. (formal talk)

CP - H.-P. Dommel and J.J. Garcia-Luna-Aceves, ``Floor Control for Activity Coordination in Networked Multimedia Applications'', Proc. 2nd Asian-Pacific Conference on Communications (APCC)'95, Osaka, Japan, June 12-16, 1995. (formal talk)

CP - H.-P. Dommel and J.J. Garcia-Luna-Aceves, ``Design Issues for Floor Control Protocols,'' Proc. IS&T/SPIE Symposium on Electronic Imaging: Multimedia Computing and Networking, Vol. 2417, San Jose, CA, Feb. 1995. (formal talk)

CP - Chane Fullmer and J.J. Garcia-Luna-Aceves, ``FAMA-PJ: A Channel Access Protocol for Wireless LANs,'' Proc. ACM Mobile Computing and Networking '95, Nov. 14-15, 1995. (formal talk)

CP - Chane Fullmer and J.J. Garcia-Luna-Aceves, ``Floor Acquisition Multiple Access for Packet-Radio Networks”,, Proc. ACM SIGCOMM 95, Cambridge, MA, August 30-September 1, 1995. (formal talk)

CP - Shree Murthy and J.J. Garcia-Luna-Aceves, ``Congestion-Oriented Shortest Multipath Routing,'' Proc. IEEE INFOCOM 1996, San Francisco, March 1996. (formal talk)

CP - Shree Murthy and J.J. Garcia-Luna-Aceves, ``A Routing Protocol for Packet Radio Networks,'' Proc. ACM Mobile Computing and Networking, Nov. 14-15, 1995. (formal talk)

CP - Shree Murthy and J.J. Garcia-Luna-Aceves, ``Dynamics of a Loop-Free Path-Finding Algorithm,'' Proc. IEEE Globecom '95, Singapore, Nov. 13-17, 1995. (formal talk)

CP - J.J. Garcia-Luna-Aceves and Shree Murthy, ``A Loop-Free Path-Finding Algorithm: Specification, Verification and Complexity,'' Proc. IEEEINFOCOM 1995, Boston, MA, April 1995. (formal talk)

CP - Shree Murthy and J.J. Garcia-Luna-Aceves, ``A More Efficient Path-Finding Algorithm,'' Proc. 28th Asilomar Conference, Pacific Grove, CA, October 31-November 2, 1994. (formal talk)

CP - Jochen Behrens and J.J. Garcia-Luna-Aceves, ``Distributed, Scalable Routing Based on Link-State Vectors,'' Proc. ACM SIGCOMM 94, London, UK, October 1994. (formal talk)

CP - Shree Murthy and J.J. Garcia-Luna-Aceves, ``A Loop-Free Algorithm based on Predecessor Information,'' Proc. ICCCN 1994, San Francisco, CA, 1994. (formal talk)

C - G. Langdon, R. Antonucci, Skip Macy, and Dean Long, "On Compression of Data from a Meteorological Stations", Proc. Data Compression Conference, IEEE Computer Society, Snowbird UT, p 471, 1994. (Poster Session)

C - J. Spring and G. Langdon, "A Study of Edge-based Lossy Image Compression Algorithms", Proc SPIE, vol 2564, Applications of Digital Image Processing XVIII, San Diego CA, pp 2-10, July 1995. (Formal Talk.)

AWARDS / RECOGNITION

Our initial REINAS work on uncertainty glyphs and data quality issues when interpolating sparse data has been leveraged to obtain NSF funding for a broader work in this area. Also received software grants of Insoft Open DVE, software and manuals, for collaborative software development to support REINAS visualization research, 1995, $30,000.
Funding for research on wireless networks has been obtained from ARPA, and will use the REINAS context as a tested.

WORLD WIDE WEB PAGE URLs:

- REINAS: http://csl.cse.ucsc.edu/reinas.html
- Environmental Visualization: http://www.cse.ucsc.edu/research/slvg/envis.html
- Uncertainty Visualization: http://www.cse.ucsc.edu/research/slvg/unvis.html
- SlugVideo: http:/sapphire.cse.ucsc.edu/SlugVideo/dream-inn.research/slvg/envis.html

[keywords{Real-time, System Design, Environmental, Sensor, Data Management]

Data Compression of Scientific Real-time End-use Data for Remotely-connected Interactive Visualization Workstations

AASERT grant number Award #:N00014-92-J-1807
Dates: 5/1/93-4/30/95 with a no-cost extension to 4/30/97

Scientific Objective:

The objective is to investigate ways to compress data going out of the REINAS database on the way to the end-user. The main interest is the interactive scientific visualization application.

Approach:

The approach in data compression is to obtain test images or test data, and investigate how well algorithms perform on the data. Work was also done on ``memory-to-memory" compression algorithms that can be called as a function by a server in a pipeline fashion.

Scientific Results:

Between September 1, 1994 and August 31, 1995 (no students were supported under the AASERT in Summer 95), we investigated time-series containers and a VQ/BTC hybrid scheme for the lossy compression of images.

Data compression on the way out of the data base is expected to by time-series containers of data from particular instruments over a given period of time. The data is compressed losslessly for transmission to the end-user. Devdutt Sheth worked on this problem in the Fall of 1994, studying the prediction error approach, and developing an adaptive compression algorithm for the wind speed and direction.

An undergraduate, Jim Spring, was supported to investigate the lossy compression of still images representing the result of a scientific visualization algorithm for presenting data from the REINAS database. Jim studied ways to adapt the VPIC (Visual Pattern Image Compression) algorithm to this application. The VPIC algorithm is a hybrid of vector quantization (VQ) and block truncations coding (BTC) that was developed by Prof Bovik and students at UT Austin was investigated for this application. The attraction was simplicity based on using a standard set of visual patterns.

Several variations were developed, and reported in a conference publication and formal talk.

Related work:

Since the Fall of 1993, under funding from the REINAS grant and this AASERT, William Macy has studied lossy methods for still image compression. These include JPEG and VQ. He also implemented a very simple lossless algorithm called FELICS (Fast Efficient Lossless Image Compression System) that gets competitive compression very simply. This algorithm is attractive for workstations that are not capable of rendering an image. In the case where the workstation is an SGI machine. Macy investigated the following methods for compression the Graphics Primitives: LZW, LZ77, JPEG. and gif. Macy also created a version of gzip that could be called from inside a server program to do the compression on the way out of the database, and reverse the process on the way to the screen of a workstation.

Publications:

Jim Spring and Glen Langdon, "A study of Edge-based Lossy Image Compression Algorithms", Proc SPIE, vol 2564, Applications of Digital Image Processing XVIII, San Diego CA, pp 2-10, July 1995.

Integration of Heterogeneous Real-time Data Repositories for Scientific Use

AASERT Award #: N00014-93-1-1038 Dates: 9/1/93-5/31/96

Scientific Objective: The original REINAS design proposal described a simple heterogeneous database system. Each organization could maintain it's own copy of the REINAS database if it wished, and each of these ``component'' databases would be networked together to form a single, logically consistent entity called a logical database network. To make a system such as the REINAS database network useful (especially to users without prior experience with relational databases), one need which must be addressed concerns query processing. The problem is this: The method of planning and processing a query that accesses multiple component data- bases that is most intuitive to a user may actually be very inefficient. Users should not need to deal with planning queries between separate databases themselves. User queries should be transformed into more efficient equivalent forms and processed automatically.

Approach: The development of efficient query processing in multidatabase management systems (MDBMS) requires solutions to many problems in query processing. A substantial amount of research has been published in this area. Much of this work has been directed toward heuristic optimization algorithms. While these algorithms provide practical engineering solutions, they frequently offer little framework for evaluating the quality of their outputs and little understanding of their effect on the performance of the MDBMS under different workloads. As a result, the performance of a MDBMS can vary unpredictably.

We have attempted to develop a theoretical model for operation allocation and operation ordering in query processing. We will use this model to find approximation algorithms whose effect on response time and processing cost can be more accurately described. The use of these algorithms in MDBMS will result more predictable response times for different workloads.

Scientific Results: Between September 1994 and September 1995, we developed a graph theoretic model for query execution. This model uses an information size metric to characterize the problem of locating the query execution with the least cost as a Single-Source Minimum Path Problem. Queries can be optimized in this model by using relational join operations as reducers. We found that while the problem of enumerating query strategies for this problem is computationally impractical in general, it can be solved in polynomial time for chain queries.

Related Work: Many heuristic solutions to allocation and ordering problems in query processing are known to exist. Recent work has focused on a variety of heuristic approaches and for a variety of different environments. There are comprehensive surveys of research in distributed query optimization.
The information size metric was chosen because it is a quantity related to both response time and processing cost: the amount of information which must be processed to complete a query. The basic assumption that makes this quantity a optimization metric is this: by reducing the amount of data that the query must process, the amount of time spent manipulating data in main memory (CPU costs), the amount time spent reading from and writing to disk (I/O costs), and the amount of time spent sending data over the network between sites (network costs) will also decrease. The smaller CPU costs, I/O costs, and network costs will reduce both processing cost and the response time.

Work by Orlowska and Zhang presented the information size metric for examining the minimizing transmission cost during query processing. They also showed how this metric can be incorporated in a query execution model to reduce overall communication cost by characterizing the problem as an Integer Linear Program.

Publications: This work has been submitted to the 1995 Annual ACM Conference for the Special Interest Group on the Management of Data (SIGMOD).

Collaborative Visualization Environments

AASERT Award #: N00014-94-1-0688

In this project, we are designing, analyzing, and implementing computer communication protocols that can support collaborative distributed applications, with the ultimate goal of supporting collaborative visualization applications being developed as part of the Real-time Environmental Information Network and Analysis System (REINAS), which is a distributed database environment supporting both real-time and retrospective regional scale environmental science. Continuous real-time data is acquired from dispersed sensors and input to a logically integrated but physically distributed database. An integrated problem-solving environment supports visualization and modeling by users requiring insight into historical, current, and predicted oceanographic and meteorological conditions. REINAS supports both collaborative and single-user scientific work in a distributed environment.

The **research problems** addressed in our research are:

- Defining a semantics for collaborative actions and events.
- Realizing floor control for coordination of collaborative activities.
- Merging floor control with a hierarchical model of session management.
- Developing a seamless zoomable graphical **group-interface**.
- Exchanging multimedia data efficiently and reliably over long-haul networks that include the Internet and radio links.
- Over the past year, we have developed new protocols and techniques for
- Floor Control Protocols: We are developing protocols and techniques to coordinate concurrent access to shared resources in real time.
- Multimedia Transport over Wireless Networks: We are developing new channel access protocols to permit the transfer of multimedia information over radio links; an example of the use of such protocols is in the distribution of live video from remote sensors to end users.
- Reliable Multicasting: We are developing protocols for the reliable distribution of multimedia information from multiple sources to multiple destinations.

Instant Infrastructure and Distributed Resource Management

Award #: N00014-95-1-1290 Dates: 8/1/95-6/30/98

Currently most mobile computing work is done within the confines of a predefined support network but, rapid deploy-ability and flexibility make mobile computing well suited to provide its own infrastructure and extend beyond these limits. Additionally, such infrastructure dependent models are ill-suited to support mobile computing in remote regions. The need to extend mobile computing to remote regions with no existing network infrastructure requires REINAS to develop techniques that will allow mobile computers to provide their own support network. We have proposed investigating mobile systems that would make greater use of the information available to them in order to construct a rapidly deploy able, highly dynamic, cost-effective instant infrastructure. The nature of mobile links is that bandwidth is limited and latency may be much higher than wired links.

Distributed environments where each node contributes to the total infrastructure have already proven themselves useful. In the environment we have proposed we assume that nodes are geographically distributed, share a common channel of communication and have a relatively large storage capacity. Transmission range and power limitations restrict the connectivity of a node to some subset of other nodes. Since nodes are mobile it is possible for them to move about, dynamically changing the region of network coverage. As they connect with each other, nodes exchange service requests. These requests propagate throughout the network until they reach a node that is able to service them. Although nodes share a communications channel and large local storage, they may be otherwise completely dissimilar.

While the REINAS project focuses on meteorological and oceanographic instruments, our mobile computing research is applicable to many other fields such as: a mobile sales force within a metropolitan region, space exploration, a battle group of ships (possibly integrating air and land support services), a disaster response team, or a firefighting team. As we extend the coverage of REINAS we are faced with a situation where mobile computing has stressed the limits of existing infrastructures. We believe that mobile computers are capable of providing their own infrastructure. By basing this infrastructure on epidemic replication we intend to realize an intrinsically reliable, cost efficient instant infrastructure that ill be rapidly deploy able, provide the flexibility that mobile computing requires and enable disconnected operation. To realize this network we must face challenges such as providing trust between nodes, how to effectively exploit varying characteristics, caching, latency reduction, efficient use of the limit bandwidth, supporting disconnected operation, and distributed resource allocation.