REINAS: REAL-TIME ENVIRONMENTAL INFORMATION NETWORK AND ANALYSIS SYSTEM -- Annual Report 1997

Principal Investigator: Patrick E. Mantey, Ph.D. Jack Baskin Professor of Computer EngineeringJack Baskin School of Engineering
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.
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
Monterey Bay Aquarium Research Institute
Pacific Grove, California
email: grbr@mbari.org
telephone: 408-647-3733

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

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

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

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

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

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **30 SEP 1997**	2. REPORT TYPE	3. DATES COVERED **00-00-1997 to 00-00-1997**

4. TITLE AND SUBTITLE **REINAS: Real-time Environmental Information Network and Analysis System --Annual Report 1997**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **University of California, Santa Cruz,Jack Baskin School of Engineering,Santa Cruz,CA,95064**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a REPORT **unclassified**	b ABSTRACT **unclassified**	c THIS PAGE **unclassified**	**Same as Report (SAR)**	**16**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

LONG TERM GOALS:

The ultimate 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. 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 initiative is to develop the capability for real-time collection, assimilation, and display of in-situ and/or remotely sensed environmental data via research in methodologies for data collection, compaction and storage as well as appropriate physical analysis and modeling. Integration of database 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.

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 provides the basic 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 a system created to support the required data acquisition, data management, and visualization, enabling real-time utilization of advanced instrumentation in environmental science, and to deploy this systems as an operational prototype. Specific technical objectives for the REINAS system are that it:

- Support interactive user access to real-time and retrospective data;
- Provide fast response for 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 to avoid data from communication link failures;
- 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 faculty and students in computer engineering and computer science, in cooperation with environmental scientists. The REINAS system has been designed for regional real-time monitoring and analysis, 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).
- A set of tools to support visualization of data, directly from the database.

REINAS user groups include: 1) operational forecasters who monitor current conditions, view standard data products, synthesize new data views, and issue forecasts and warnings; 2) modelers who analyze new model products and compare them with other models and with past and present conditions; 3) experimental scientists who 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; and 4) instrument engineers who add new equipment to the system, access and update metadata describing individual devices and methods of calibration, study maintenance records, and profile sensor quality. 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 both new and existing instruments belonging to the participating institutions, plus those of the National Weather Service and the Monterey Bay Aquarium. A collection of standard load paths is used to move data from the instruments into the system. 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 perform any data parsing that is required during the first steps of the data load process. One or more commercial (relational) database systems are used as the central components within REINAS, augmented with software implementing memory-resident caches and disk resident transaction logs to buffer heavy data streams and provide reliability guarantees in case of system or network failures. The use of a relational database with a user-accessible API in REINAS provides a common data access model for arbitrary database products. Thus, researchers who implement a scientific application using the REINAS API effectively have direct access to all REINAS data in a common binary format. A set of control applications is used to deploy, configure, and steer collection devices, and a 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 coupled directly (via an API) to the database.

WORK COMPLETED :

The REINAS database design supports measurement data in both temporal and spatial organizations, plus support for metadata. Data collected by the real-time REINAS instruments are received by the REINAS/PC at the instrument site, and parsed by REINAS instrumentation software which formats and logs the data streams appropriately before transporting the data to database nodes via (sometimes unreliable) network connections. Psuedo real-time or historical data not received in real-time may arrive from a number of different paths (or ``virtual" instruments), such as archival files, Internet based File Transfer Protocol (FTP), UNIDATA broadcasts, etc. In general, any instrument that produces a consistent or predictable digital data stream can be connected to the

REINAS system.
Early in 1996, a revised third-generation architecture for REINAS was defined and developed based upon the lessons and experiences of the previous three years. This architecture is comprised of three software realms:

- instrumentation -- reliably parsing and forwarding measurement data in real-time
- system -- storing and retrieving data using relational database constructs
- end-user -- visualization, models, and related applications

This architecture has proven to be extremely successful and flexible. One of the most important aspects of this new approach is the REINAS ``dispatcher'' -- the system-node query-interface which any and all REINAS client applications use to query and interact with the system. The dispatcher includes the definition of a flexible but easily learned interface, which can be used interactively; much of the subsequent success of REINAS as a useful enterprise is a result of the development of the REINAS dispatcher definition. A REINAS Systems Software Manual was prepared, containing details for others wishing to install and use REINAS, including guides for REINAS installation and internals, for REINAS core system APIs, and for writing REINAS device managers. In January 1997, to demonstrate the efficacy of the REINAS approach to environmental monitoring in general as well as the ease and portability of the software solutions, REINAS was ``ported'' to a separate regional environmental monitoring effort centered on the Columbia River estuary near Portland, Oregon. A two-day seminar and introduction to REINAS was held at the Oregon Graduate Institute by REINAS personnel. Shortly thereafter, OGI scientists and engineers deployed REINAS software to solve their own local data collection and distribution needs.

A set of CODAR SeaSondes (high-frequency Doppler radars manufactured by CODAR Ocean Sensors) connected to REINAS in real-time provide measurements of radial ocean surface currents from three fixed positions surrounding the Monterey Bay, typically at half-hour or hourly intervals. Using software developed at UCSC, radial data from two or more of these CODAR instruments automatically generate vector maps of surface currents at resolutions of approximately 3 km; along with the radial data, these derived ocean surface vectors are available within REINAS. In the past year, the National Weather Service Next-Generation Weather Surveillance Radar (WSR-88D) or ``NEXRAD'' located atop Mount Umunhum and producing high resolution three-dimensional Doppler datasets of the atmosphere above Monterey Bay, the Salinas Valley, and the San Francisco Bay region, was integrated, into REINAS. The NEXRAD produces extremely large (multi-megabyte) datasets which pose significant challenges to any application seeking to distribute, process, or archive the data sets in real-time. REINAS currently supports receipt, storage, and distribution of NEXRAD data at a very fine granularity.

RESULTS:

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.

REINAS has implemented a regional wireless, real-time network based on the Internet Protocol (IP) for remote data collection within the Monterey Bay and surrounding area. The network currently comprises over 3000 square miles and spans three counties, the Monterey Bay, and immediately adjacent areas of the Pacific Ocean. Within this network, twelve surface weather stations, three Doppler ocean-current radars, two video cameras, one submerged sea temperature and salinity instrument, one vertical wind-profiler radar, one next-generation Doppler weather surveillance radar (NEXRAD), and other experimental radars feed real-time data to the REINAS system. Additional instrumentation, including twenty-seven additional surface weather stations, two ocean surface weather stations, two submerged ocean thermistor chains and one acoustic Doppler ocean current profiler provide near real-time data to REINAS using the conventional Internet. Recent deployments of other 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.

In previous years, investigations into the development and implementation of remote video streams as real-time REINAS instruments were initiated. At a system level, this included defining a prototype video instrument, building or acquiring the necessary hardware and software, deploying the camera and 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 later that spring. The high data-rates achievable with a remote camera instrument (64 kbps to 8 Mbps depending on frame rate, content, and other factors) motivated a research effort into the issues surrounding supporting robust transmission of real-time video streams over lossy (wireless) digital networks. In addition, because the camera platform is remotely maneuverable, ongoing research efforts into producing near real-time seamless high-quality panoramic imagery of the sea-state and atmosphere are also underway.

Since May 1996, the third-generation of REINAS system software has operated online in an operational sense – supporting data collection and distribution in a continuous (24 hours/day, seven days/week) mode. During this time, REINAS has successfully operated in an operational mode, essentially without interruption, and continues to incorporate new data sources on an ongoing basis. To date, thousands of users among the general public access REINAS daily through predefined query-pages on the World Wide Web. A select set of research scientists at the Naval Postgraduate School, University of California, California State University, MBARI, Monterey Bay Aquarium, California State Pacific Fisheries, and Long Marine Lab, among others, also use REINAS as part of their own research efforts.

PUBLICATIONS

P - Eric C. Rosen, Theodore R. Haining, Darrell D. E. Long, Patrick E. Mantey. ``Mechanisms for Reliable Information Flow In REINAS". Submitted for publication, 1997.

P - Eric C. Rosen, Theodore R. Haining, Darrell D. E. Long, Patrick E. Mantey, and Craig M. Wittenbrink. ``REINAS: A Real-time system for Managing Environmental Data". Journal of Software Engineering and Knowledge Engineering, 1996.

Rosen, E. C. ``Evaluating Digital Video Stream Transmission via Packetized Wireless Channels". Ph.D. Dissertation. UCSC Technical Report UCSC-CRL-96-31, 1996.

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 database 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 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.