REINAS: A Real-time System for Managing Environmental Data †

Darrell D. E. Long, Patrick E. Mantey, Eric C. Rosen, Craig M. Wittenbrink
Baskin Center for Computer Engineering & Computer Science
University of California
Santa Cruz, CA

Bruce Gritton
Monterey Bay Aquarium Research Institute
Moss Landing, CA

Abstract

Managing scientific data is a challenging task, and many of the problems it presents have yet to be adequately solved.

The Real-time Environmental Information Network and Analysis System (REINAS) is an attempt to develop an operational solution to the problem of collecting and distributing environmental data in a real-time context, as well as supporting data acquisition, verification, retrieval, visualization and long-term data maintenance. The system is built around one or more databases and has been developed to support 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 system. In such a system, the database can provide a powerful structure to handle data management, but current database technology has difficulty meeting the performance requirements that a large real-time environmental system demands.

This discussion will describe the REINAS architecture in some detail, including challenges that were addressed in the construction of an operational system, with real users requiring real-time and retrospective access to a variety of structured environmental data, while supporting continuous real-time collection of data into the database.

†Supported by the Office of Naval Research under grant N00014-92-J-1807 and by the National Science Foundation under grant IRI-9423881.

1 Introduction

The Real-Time Environmental Information Network and Analysis System (REINAS) (described at earlier stages of development in [9, 10]) supports regional-scale environmental science, monitoring, and forecasting through a distributed system connected by the Internet. It is being developed by the University of California at Santa Cruz (UCSC), in conjunction with the Naval Postgraduate School (NPS), and the Monterey Bay Aquarium Research Institute (MBARI).

With REINAS, environmental scientists have the ability to observe, monitor, and analyze regional oceanographic and meteorological phenomena from their desk top. It has been designed to provide:

- 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
- A logically consistent distributed database that stores data independently of file format and which maintains metadata describing where and how data was obtained (*i.e.* the database tracks data pedigree).

REINAS serves different user groups: *operational forecasters* who monitor current conditions, view standard data products, synthesize new data views, and issue forecasts and warnings; *modelers* who analyze new model products and compare them with other models and with past and present conditions; *experimental scientists* who collaborate with other scientists on-line, observe individual data fields as they are collected, and who may modify data collection methods for an ongoing experiment; as well as *instrument engineers* who add new equipment to the system, access metadata describing individual devices and methods of calibration, study maintenance records, and profile sensor quality. In addition, the *general public* (especially the local community of sailors and wind surfers) form an ever-growing group that has come to depend on real-time access to REINAS environmental data via the World Wide Web [13, 21].

Interactive real-time measurement, real-time monitoring, and retrospective data management using REINAS provide forecasters and experimental scientists with a framework to plan their experiments, expeditions, and monitoring activities. The special focus on real-time observation and analysis allows the oceanographic and meteorological communities to identify phenomena as they occur and to react to emerging phenomena and trends by changing the nature and frequency of data collection at sites of interest.

The integrated problem solving and visualization environment provides scientists and forecasters with pictures of environmental features, trends, and relationships to study dynamic environmental behavior in a geographic context. Scientists can fuse data collected in the field with data generated by numerical models and simulations.

System configuration tools simplify instrument engineering tasks, for instance, supporting easy addition of new instruments to the system. Instruments are connected to REINAS by both radio and terrestrial links of various types and data rates.

The data base supports access to data, collected by different people and institutions at different times, in a logically consistent manner. The data base separates users from many of the mechanical tasks of data management, as it is designed around an architecture integrating real-time data from multiple instrument technologies, classes (numeric, text, and video), and representations. Both scientific data and metadata are stored using a stable schema which maintains data pedigree. An interactive electronic log book allows experimental scientists to keep notes easily accessible.

2 Related Work

There are several similar systems to REINAS, and we briefly review some of them. Sequoia 2000 was a large multi-year development project which focused on global environmental science [18, 19]. Sequoia 2000 differs significantly from REINAS in that while REINAS is driven by real-time constraints, Sequoia 2000 was primarily designed for retrospective access with little real-time emphasis. Like REINAS, Sequoia 2000 was built around a database [15, 17]. Real-time systems employing less sophisticated database technologies have also been developed [3, 6, 16], as well as others that run at much lower data rates [5]. One such system, StormCast, collects and archives environmental data primarily for current and forecast use, as well as supporting operating system technologies research [7]. Unlike REINAS, Stormcast does not address the long-term data management problem, the issues involved in overlaying different types of structured scientific data, or the advantages of treating instruments as active Internet nodes instead of passive data-streams. The NEONS project [12] did employ a relational database engine for environmental data, but unlike REINAS, NEONS did not emphasize real-time access to instrumentation. Some previous Geographic Information Systems employ structures similar to that used by the REINAS database schema [1, 2], while some recent data visualization work [8] is similar to that developed within REINAS. REINAS differs from these previous works in that it stresses very short time-scale real-time access to environmental data through a robust and highly structured database system, supporting regional three-dimensional visualization products as well as custom user-level applications through a generalized query interface.

3 The System

REINAS is a large, distributed, system, made up of literally hundreds of components scattered across a geographic area of several hundred square miles. The architecture of its various subsystems has evolved over the life of the project into three major pieces.

3.1 Database Nodes

REINAS includes one or more database systems, into which real-time environmental data flows from dispersed instrument nodes. When data arrives, it is handled by a merge-server, which immediately logs the data to disk. A separate db-loader reads the log and formally inserts the data into the database; shared access to the log is coordinated by a robust log-manager. In addition, the merge-server forwards the most-current sensor values as they arrive to the real-time cache, which maintains the current state of the sensor network outside of the database. Finally, queries to the system are received by the dispatcher, which queries the database or the real-time cache as appropriate.

The structure in Figure 1 reflects several operational and technological constraints. For example, at times, the input data-rate may exceed the ability of the

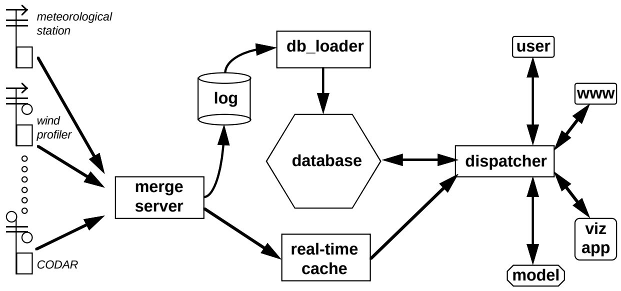

Figure 1: REINAS System Data Flow

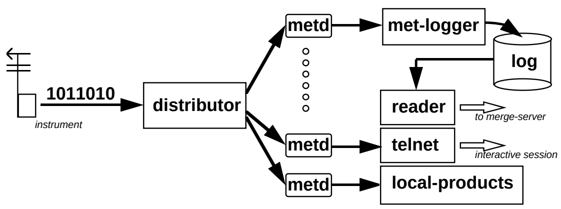

Figure 2: Instrument Node Data Flow

database to accept new data; the log provides a reliable mechanism to buffer the receipt of data from the rest of REINAS, allowing the system to continue to accept data, keep track of real-time values, and eventually migrate all data into the database. Similarly, the real-time cache exists to provide easy access to real-time sensor values, since *snapshot* queries for the most-recent data tend to dominate among the number of different types of queries received.

3.2 Instrument Nodes

One of the unique aspects of REINAS is the emphasis it places on real-time data collection, and the fact that the system was conceived and built to be tightly integrated with the Internet. As a result, instruments within REINAS are accessed directly through Internet nodes, and data is forwarded from these remote nodes to a database node of the form presented in Figure 1. In practice, this often means deploying a computer near the remote instrument and networking the computer to the Internet via wireless or land links, as appropriate.

Raw environmental data typically arrives in computer form through a serial or parallel. This stream is managed by a single process, the distributor, which performs all instrument-specific interactions, parses the stream, and implements instrument access control. Through a paired UNIX daemon, the distributor allows multiple client processes to gain shared access to the stream, and when appropriate, exclusive access. The daemon provides a separate standard interface to processes which perform actions on the instrument data itself, such as forwarding the data to a REINAS database node. Since the instrument is typically at the remote end of a slow and unreliable network link, a logger logs the stream to disk and a separate reader transmits each data record to a merge-server at the database node, deleting the record from the log only after the merge-server has acknowledged receipt.

Figure 2 illustrates this software architecture applied to the case of a generic meteorological station. A single distributor tuned to a specific flavor of meteorological instrument, such as those provided by a particular vendor, manges the stream. Multiple instantiations of the metd daemon interact with the distributor and provide a generic meteorological station interface. The met-logger logs the stream, while a localproducts process maintains a raw local archive. Scientists, system-administrators, and even specialized applications may also interact with the stream or instrument in parallel.

3.3 Network

The instrument and database nodes, as well as users running visualization applications, are connected via the Internet. In places the network is built over conventional wire links (*e.g.* Ethernets, T-1 bridges); elsewhere, it is implemented using wireless (radio) links of varying bandwidth. The characteristics and performance of the network is a crucial aspect of the system as a whole, and especially in the case of wireless networks, tuning and new wireless-aware protocols are necessary to achieve satisfactory reliability and performance.

4 The Real-Time Problem

The most unique and challenging aspect of REINAS is its real-time focus, and the fact that the time-scale of the environmental phenomena being monitored determines the real-time performance requirements.

4.1 What Does Real-time Mean?

A real-time system means different things at different levels. The user's perspective is often determined by the application interface and the type of feedback provided. A visualization application that allows the user to manipulate static data interactively might be considered real-time by some, even though the data itself may not change. Other users may define a system to be functioning in real-time only when they can perceive that the data is being updated frequently. REINAS uses a best effort network, and is therefore not real-time in a systems sense.

Under REINAS, real-time is defined at the level of the environmental phenomena being observed. For some sensors attached to surface weather stations (such as those directly measuring wind gusts), sub-minute sampling rates are appropriate and real-time in this context means updated values arriving every 10 seconds. For others sensors, such as air-temperature, 60 second sampling periods are appropriate. Radars measuring slower changing phenomena such as ocean currents tend to require longer term averaging, and real-time in this context may infer updates occurring at 30 minute or one hour intervals.

4.2 Real-time Challenges

Typical real-time database systems are oriented toward applications that heavily favor read operations over writes. Examples include airline reservation systems, automatic teller machine interactions, and point-of-sale credit verification systems.

In REINAS, separate inserts from several hundred distinct time-series sensors are independently issued every few seconds. Other more structured data arrive for insertion at slightly less frequent intervals (from every minute to every few hours hours). User-level regional-visualization applications provide the heaviest demands for reads, but the maximum level of retrievals is much smaller than the insertion rate. Note that in a real-time environmental database system such as REINAS, updates of existing data are rarely performed. Hence, system performance is constrained primarily by the capacity of the database to incorporate a heavy, steady stream of new data (typically hundreds of inserts every few seconds).

Because current commercial database products cannot keep pace with the high rate of inserts that REINAS must support, inserts migrate through a disk-based *log manager* where they are buffered into more efficient bulk loads. Each log is produced by a merge-server and consumed by a db-loader. If necessary, multiple logs can be maintained.

Commercial databases may also have problems keeping pace with the demand for data, especially real-time data. In fact, real-time access to current data represents the most popular type of query among the user community (in terms of the number of separate queries). To alleviate the demands placed on the database, a cache of real-time sensor values is maintained externally in the real-time cache; the query dispatcher retrieves real-time data directly from this cache instead of the database. Providing seamless access to all REINAS data regardless of whether it exists in the database itself or the real-time cache, is among one of the software engineering challenges facing the project.

5 Scientific Data Models

The REINAS approach to data management is dictated by the requirements imposed by all phases of environmental science. The fundamental design decision was to integrate the management of scientific data and meta-data, rather than treat these data types separately (or ignore metadata) as is common in other scientific systems. Hence, REINAS developed an information model to structure the activities, interactions, and products relevant to the users and processes performed from experiment planning to data archiving, as well as a structure for the various types of sensor values. The model integrates the specification of *scientific data* such as measurements, sensor values, and other observations, with the specification of *scientific metadata*, including definitions for the content, representation, structure, and context of the scientific data.

Scientific data in general, and data streams in REINAS, can be broadly categorized into one of two distinct classes: environmental or *sensor-produced* data streams, and meta-data or *data-about-the-data*. REINAS was designed to handle data of each type, as well as many relationships between them.

5.1 Environmental Data Streams

In its early stages, REINAS focused upon managing data from three specific types of instruments: surface meteorological (MET) stations, Coastal Ocean Dynamics Application Radars (CODARs), and vertical wind-profilers. These core instruments provided a rich set of environmental data-streams; MET stations generate multiple parallel streams of time-series data (scalar and vector valued), profilers produce vertical arrays of vector and related scalar data, while CODARs generate two-dimensional arrays of vectors organized either radially or, when combined from multiple sites, on a regular grid bounded by the local coastline.

The simplest type of environmental data is a sensor value, which may be comprised of an aggregate of several scalar values. When sampled temporally, the resulting *time-series* data represents the simplest data-type an environmental database system must be able to manipulate (Figure 3). REINAS treats time-series data as a fundamental data-type and handles a variety of different flavors; the most common are *meteorological data* from surface weather stations. Most meteorological sensors, such as those measuring air-temperature, humidity, or solar-irradiance, present a time-varying scalar

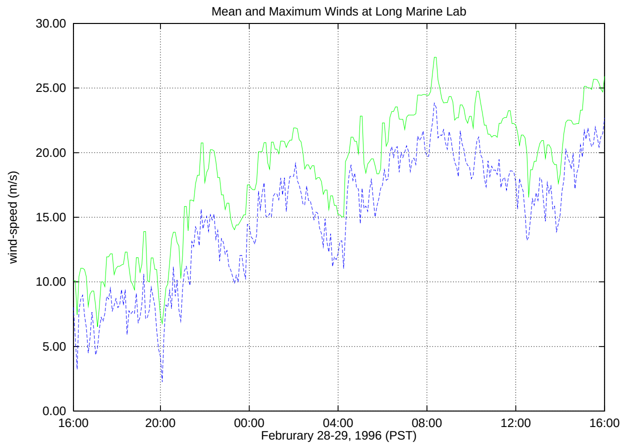

Figure 3: Sample Time-series data

quantity. Others, such as a wind-monitor, present a vector, *i.e.* ⟨*wind-speed, wind-direction*⟩.

Doppler radar technology allows *wind-monitor*-type data to be indirectly measured in a vertical column of the atmosphere, at intervals of few hundreds of meters to heights of several thousand meters. *Wind-profilers* thus produce a vertical profile of a tall column of air; an acoustic doppler current profiler can perform the analogous measurement using sound for a water column. Hence, vertical profiles of wind or water currents, possibly augmented with other indirectly measured variables, and updated routinely over time, are also a basic environmental data stream. Such profiles are derived from radar-return spectra, which may also be included in the data base.

Another type of radar, the CODAR, indirectly measures ocean-currents along a collection of radials from the radar site at fixed range intervals to distances of a few tens of kilometers. When *CODAR radial data* from two or more appropriately positioned CODAR sites are combined, *a field of ocean surface-current vectors* can be generated (Figure 4). Each radial or vector may also be accompanied with statistical confidence measures. A set of CODAR vector maps collected over time can be used to produce an animation illustrating ocean surface-current trends.

These three data sources require REINAS to manipulate scalar and vector data in point, linear, gridded, and higher-dimensional structures. Many other environmental data streams are similar in construction to one of the core types and can be incorporated fairly easily. REINAS has already been extended to include other interesting environmental data-streams, including multi-spectral satellite imagery, rawinsondes, and three-dimensional regional and global model outputs. In addition, work in progress will extend the system to

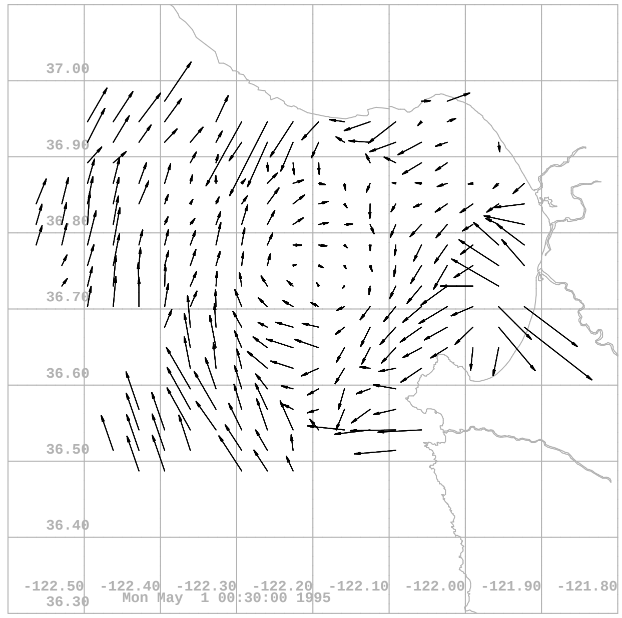

Figure 4: Sample CODAR Vector Data

include video streams from remote camera instruments (the prototype of which is already a popular feature on the World Wide Web [14]).

5.2 Schema

In order to implement a data model to support a growing, and unknown set of future instruments, a schema was designed that allows new instruments to be added to the system without requiring the creation of new tables. Only when data-streams of a type not previously supported are encountered must the schema be augmented. The modeling and schema development process was started with the MBARI Scientific Information Model (MSIM). MSIM was developed to establish a generalized framework for the long term management of interdisciplinary ocean science data. MSIM's requirements included the need to capture information associated with data collection activities and systems, processing algorithms, collection and processing of laboratory samples, quality assessments, scientists and technicians, experiment sites, and sampling protocols. Although MSIM provided a good starting point, REINAS required a richer description of some metadata concepts. Specifically, direct connection and control of instrument systems required a richer description of current and historical configuration of platforms and instruments. In addition, since instrument systems would belong to multiple institutions in a *regional laboratory* context, REINAS required the ability to store scientific observations of the same environmental parameter with multiple logical and physical representations.

Further, to accomplish full data integration the REINAS model must support multiple data classes (numbers, sets of numbers, text, images, video, sound), structured versus unstructured data, multiple aggregation schemes for observations (point data, profile data, two and three-dimensional fields), and multiple sources (in-situ, remotely sensed, mobile platforms, historical files and databases).

Given this complex set of requirements, REINAS made two simplifying assumptions. First, it was assumed that some level of conceptual generalization applies across all disciplines of the earth sciences, including those most relevant to REINAS: meteorology and oceanography. For example, the concepts of platforms (*e.g.* ships, buoys, airplanes, satellites), instruments (*e.g.* weather stations, profilers, CODARs), and sensors (*e.g.* temperature, humidity, wind) can be described by a common model of *systems* with appropriate subtype specifications. Second, it was assumed that the primary access would be by type of data or environmental parameter. The physical layout of the database and schema provide the fusion of all environmental parameters from all instruments. Enough relational links are provided to track back to the collection instrument, and the collection circumstances.

From these simplifying assumptions, the REINAS database is defined from a single, self describing schema to which load paths from all sources are directed. Standard templates are used for developing and reusing load path software from related sources. Since this template links with the schema capturing this metadata is performed while capturing the primary observation data. More importantly, all information required for the access and assessing of data from multiple sources is available. The data model provides meaningful selection criteria, supports the translation of logical concepts, supports the exchange of data stored in differing physical formats, and supports long term data management.

6 Output Products

The data residing within REINAS can be retrieved and used in a variety of ways. Although queries for current-conditions tend to dominate, fairly complex queries are also possible. For some purposes, a *Structured Query Language* (SQL) [4] interface for submitting queries is sufficient. In general, however, users and user-level applications require a more abstract interface for requesting and receiving data from REINAS.

6.1 Query Types

Regional environmental data may be queried in a variety of ways. The most numerous type of query issued is the request of current or real-time sensor values. Often these queries pertain to a specific instrument site and hence they have been labelled *snapshot queries*. These queries return one or more sensor values.

It is also possible to query over a broader set of environmental data. For example, one can ask for a plot of the previous hour's mean ocean surface current speeds at every CODAR vector grid-point when the air-temperature at the Monterey Bay Aquarium dropped below a previously recorded minimum air-temperature for the region. Such complex queries may involve manipulating and returning large amounts of scientific data.

REINAS queries may involve querying non-scientific data, and the ability to support such queries has been one of the main goals of the project. Being able to query for the ten year old calibration of a particular temperature probe may be essential for a researcher studying long term trends. Such queries return data about the collection process itself, which helps in the assessment of the validity or uncertainty of the data, and also makes it possible to compare data collected in different times, with different processes, and with different instruments.

6.2 API

Data residing in a well-designed, well-maintained database has no value if it cannot be retrieved for use by user-level or other applications. The scope of useful applications is too broad to allow a system to include all of them; useful applications are always being developed. As a result, users should be able to develop their own applications through a predefined interface.

Initially, only one such *application-programmer-interface* existed within REINAS. A low level *REINAS-API* was constructed essentially to provide a generic framework for issuing SQL-level operations against the database. It was quickly realized, however, that only a fairly specialized group of individuals could effectively use this API; its primary drawback is that the application programmer must understand not only SQL but also the REINAS schema, a requirement that is unreasonable in an application where the typical applications-programmer are environmental scientists with little knowledge of database technology.

Hence, a higher level API that buffers the applications-programmer from the REINAS schema and SQL semantics was designed. *RSObjects* is an object oriented API implemented within the C programming language that is being developed to provide such an interface to C and FORTRAN programmers [21]. RSObjects was developed following several other API development efforts with prototype systems. These APIs were difficult to extend, failed to isolate client applications from changes within the dispatcher and database, and were unable to support large objects (among other failures). In order to mitigate these problems, RSObjects provides a dialogue model of interaction where the application asks the database what it may request. In this way the applications maintain backwards compatibility, while also providing opportunity for querying future instruments and views without recompilation.

6.3 User Level Applications

To date, REINAS has developed many retrieval applications: generic monitoring tools using conventional glyphs and time-series plots, contour, station plots, and overlaid satellite data, and three-dimensional isosurface, streamlines, glyphs, and bump map renderings.

The three-dimensional rendered visualizations have provided a new insight into various aspects of the environmental data being archived in REINAS, and have also been the stimulus for several novel visualization methods [11, 20]. But, these visualizations are not used by the majority of REINAS users, as few REINAS users have direct access to Silicon Graphics or other accelerated rendering platforms However, many users can access with acceptable performance the graphics presented on the World Wide Web (WWW) REINAS pages. To date, via the WWW, REINAS has developed a following of thousands of users worldwide, many of whom have come to depend on the simple time-series and vector plots that provide real-time regional meteorological and oceanographic conditions [13]. Over a million separate visits have been recorded to the prototype video platform interface [14] which also displays real-time conditions from a nearby REINAS weather station. These web pages support many casual users, but also provide an effective mechanism for data access for many of the collaborators, scientists, and professional meteorologists using REINAS.

7 Conclusions

The initial REINAS proposal written in 1991 described a system architecture based on an assessment of available computer technology and a forecast of developments to be expected in the following years. The proposal included the concept of extending instrument sites into Internet nodes using a local 80x86 class computer or a "REINAS/PC", providing local parsing, processing, and buffering. Now implemented and revised, the REINAS/PC concept has been demonstrated to be a very viable and powerful architecture, providing modularity, extensibility, and operational flexibility. The Internet Protocol has already been demonstrated to be a robust and flexible foundation for building a distributed regional network using a variety of technologies (*e.g.* wireless, serial, Ethernet, ATM), while local buffering and archiving provide insurance against data loss, even in the event of extended network (or remote database) failures. Since this architecture was first developed, microprocessor and related technology costs have declined significantly, and the development of an even cheaper, single-board REINAS/PC with embedded networking options, digital signal processing and input/output capability, and internal static storage is a promising avenue for future research.

Two closely related and key decisions in the development of REINAS were the use of commercially available relational-database systems to support data capture, retrieval and long-term management, as well as the decision to use a *single* database system to support both real-time and retrospective uses of the data. Results from the use and development of REINAS to date suggest that this may have been a correct decision, in that REINAS has successfully employed a commercial relational database management system. However, some performance problems remain in acceptably supporting the growing demand from real-time users.

The real-time cache and db-loader log have overcome the "real-time problem" of data input, supporting an insertion rate driven by continuous data feeds from the instruments, while making current measurements available in real-time. Query support and the RSObjects API provide structured access to all data in the database. With the growing use of the REINAS system from remote Internet users via the World Wide Web, most of the queries focus on real-time and recent data. This heavy real-time interest was not anticipated in the original REINAS proposal (which predicted instead a set of primarily scientific users with different, longer-term access requirements) and current research is emphasizing improvements in the real-time query performance to better support the large number of ad-hoc users of the system. The challenge here is to overcome the inherent conflict in providing acceptable performance to interactive users with a database schema developed to support scientific users engaged in retrospective analysis. Additional indexing of the data, restructuring of some tables, and other database tuning efforts are being investigated to improve interactive query performance.

References

[1] T. Bernath. Distributed GIS visualization system. In *Proceedings of GIS/LIS, Vol. 1*, pages 51–58, San Jose, CA, Nov. 1992. American Society for Photogrammetry and Remote Sensing.

[2] A. Berrill and G. Moon. An object oriented approach to an integrated GIS system. In *Proceedings of GIS/LIS, Vol. 1*, pages 59–63, San Jose, CA, Nov. 1992. American Society for Photogrammetry and Remote Sensing.

[3] R. S. Cerveny et al. Development of a real-time interactive storm-monitoring program in Phonenix, Arizona. *Bulletin of the American Meteorological Society*, 73(6):773–779, June 1992.

[4] C. J. Date. *A guide to the SQL Standard: a user's guide to the standard relational language SQL*. Addison-Wesley, Reading, Massachusetts, third edition, 1993.

[5] S. Howes. Use of satellite and radar images in operational precipitation nowcasting. *Journal of the British Interplanetary Society*, 41(10):455–460, Oct. 1988.

[6] J. Intriery, C. Little, W. Shaw, R. Banta, P. Durkee, and R. Hardesty. The land/sea breeze experiment (LASBEX). *Bulletin of the American Meteorological Society*, 71(5):656, May 1990.

[7] D. Johansen. StormCast: Yet another exercise in distributed computing. *Distributed Open Systems in Perspective*, Feb. 1993.

[8] J. P. Lee and G. G. Grinstein, editors. *Database Issues for Data Visualization, IEEE Visualization '93 Workshop*. Springer–Verlag, 1994.

[9] D. Long, P. Mantey, C. M. Wittenbrink, T. Haining, and B. Montague. REINAS the real-time environmental information network and analysis system. In *Proceedings of COMPCON*, pages 482–487, San Francisco, CA, Mar. 1995. IEEE.

[10] P. Mantey et al. REINAS: Real-time environmental information network and analysis system: Phase IV - experimentation. Technical Report UCSC-CRL-94-43, Computer Engineering and Computer Science, University of California, Santa Cruz, 1994.

[11] A. Pang. Spray rendering. *IEEE Computer Graphics and Applications*, 14(5):57 – 63, 1994.

[12] C. Ramanathan. Providing object-oriented access to a relational database. In *Proceedings of the 32nd Annual Southeast Conference*, pages 162–165, New York, NY, Mar. 1994. ACM.

[13] E. C. Rosen. REINAS: Real-time environmental information network and analysis system: Home page. URL: http://csl.cse.ucsc.edu/reinas.html, 1995. Computer Engineering, University of California, Santa Cruz.

[14] E. C. Rosen. *SlugVideo* home page. URL: http://sapphire.cse.ucsc.edu/SlugVideo/, 1995. Computer Engineering, University of California, Santa Cruz.

[15] L. Rowe and M. Stonebraker. The Postgres data model. In *Proceedings of the 13th Conference on Very Large Data Bases*, pages 83–96, Brighton, England, Sept. 1987.

[16] D. Schwab and K. Bedford. Initial implementation of the great lakes forecasting system: A real-time system for predicting lake circulation and thermal structure. *Water Poll. Res. J.*, 29(2/3):203–220, 1994.

[17] M. Stonebraker. The design of the Postgres storage system. In *Proceedings of the 13th Conference on Very Large Data Bases*, pages 289–300, Brighton, England, Sept. 1987.

[18] M. Stonebraker. An overview of the SEQUOIA 2000 project. In *Digest of Papers: COMPCON Spring 1992, 37th IEEE Computer Society International Conference*. IEEE Computer Society Press, Feb. 1992.

[19] M. Stonebraker. Sequoia 2000: A reflection on the first three years. *IEEE Computational Science and Engineering*, 1(4):63–72, Winter 1994.

[20] C. M. Wittenbrink, G. Fernandez-Ubiergo, and J. Glen Langdon. Feature extraction of clouds from GOES satellite data for integrated model measurement visualization. In *IS&T/SPIE Symposium on Electronic Imaging: Image and Video Processing 1996*, San Jose, CA, Jan. 1996.

[21] C. M. Wittenbrink, E. Rosen, A. Pang, S. Lodha, and P. Mantey. Realtime database support for environmental visualization. In *Second Workshop on Database Issues for Data Visualization*, Atlanta, GA, Oct. 1995. IEEE.