

MICROCOMPUTER APPLICATIONS

VOLUME 17, No. 3, 1998

• Expert System Shell Selection for Moving Towards an Autonomous Ground Control System for the Microwave Anisotropy Probe Satellite	J. Liebowitz, S. Coyle, J. Breed, R. Hollenhorst, J. Dowling, J. Condrón, J. Stuart, J. Klein, S. Shulman	57
• Semantic Geographic Information System	E. Alvarez, N. Rische, D. Barton	61
• Multithreading Implementations	K.M.Kavi	70
• Computer Simulation in Hospital Management	L. Moreno, R.M. Aguilar, C.A. Martín J.D. Piñeiro, J.F. Sigut, A.F. Hamilton J.I. Estévez, J.L. Sánchez, V.I. Jiménez	85
Author Index 1998		93
Journals / Editors		94

SEMANTIC GEOGRAPHIC INFORMATION SYSTEM

E. Alvarez, N. Rishe, and D. Barton*

Abstract

This research describes the design and implementation of a semantic Geographic Information System (GIS) as well as the creation of its spatial database. This database schema was designed and created, and all textual and spatial data was loaded into the database, using the semantic Database Management System's (DBMS) binary database interface that is currently being developed at Florida International University's (FIU) High Performance Database Research Center (HPDRC). A friendly graphical user interface has been created together with the main system's displaying process, data animation, and data retrieval. All these components are tightly integrated to form a novel and practical semantic GIS that facilitates the interpretation, manipulation, analysis, and display of spatial data such as ocean temperature, ozone (TOMS), and simulated (Sea-viewing Wide Field-of view Sensor) SeaWiFS (SeaStar) data. In addition, this system has played a major role in the testing of the HPDRC's efficient and high performance parallel semantic DBMS.

Key Words

Database, spatial data, GIS, remote sensing, semantic database

1. Introduction

There has recently been a remarkable surge in academic, political, commercial, and scientific interest in the use of Geographic Information Systems (GIS). Many of these initiatives have been stimulated by current developments in remote-sensing. This, in turn, has created a need for the development of a more efficient and user friendly GIS.

The High Performance Database Research Center (HPDRC) [1] has also experienced this need for an efficient GIS. At the HPDRC, there is wide variety of spatial data sets from several sources including ocean temperature data (supplied by the University of Miami Rosenstiel School of Marine and Atmospheric Science), simulated SeaWiFS (Sea-viewing Wide Field-of view Sensor) (deployed by NASA from the SeaStar satellite), and ozone data (TOMS, Total Ozone Mapping Spectrometer), (deployed by NASA's Goddard Space Flight Center). Due to the large amount of data inherent in these types of data products, we found a need for a computer-based system able to efficiently store, manipulate, analyze, and display this information. At the same time, there was a need for a system to test the high performance and efficient parallel

semantic DBMS currently under development at the HPDRC. Hence, a semantic GIS provided a solution for both needs.

The purpose of this research is to design and implement a semantic GIS that uses a semantic DBMS for its spatial data storage and retrieval. Most commercially available GIS packages use a relational DBMS. These relational DBMSs, however, are lacking in a number of important areas more fully attended to by the semantic DBMS. The semantic DBMS will provide protection and security for the data and, at the same time, enforce consistency of the stored data. It will also provide efficient, simultaneous retrieval of massive amounts of data to multiple users and ensure better logical properties such as a comprehensive enforcement of integrity constraints, greater flexibility, and substantially shorter application programs [2].

The semantic GIS is designed to facilitate the displaying and interpretation of spatial data sets by providing an interface between users and the semantic database that contains the data. This application accepts run-time parameters from users and utilizes these parameters to make queries to the database and retrieve specific details about the data set under study. These details include instrument, satellite, date, frequency, data images, and many other types of information.

In sum, by building a GIS using a high performance semantic DBMS, efficient storage, manipulation, analysis, representation, and retrieval of both conventional and spatial data sets is provided. In addition, this is all implemented using a very friendly graphic user interface and a good set of analysis tools. These features are the main factors for the success of this system.

2. Background

2.1 What is a Geographic Information System?

GIS are a combination of computer hardware, software, and procedures designed to support the capturing, management, analysis, Modelling and display of geographically referenced data [3]. GIS is a sophisticated computed based mapping and information retrieval system which consists of three primary components:

- A powerful computer graphic program
- One or more external databases
- A set of analysis tools

A GIS must have all three components tightly integrated in order to be a good computer system.

* High-Performance Database Research Center, School of Computer Science, Florida International University, University Park, Miami, FL 33199, USA; e-mail: hpdrcc@cs.fiu.edu, <http://hpdrcc.cs.fiu.edu>

2.2 Features of the Semantic DBMS used in the GIS's Implementation

GISs demand high performance and pose some very special requirements for database management. These systems require special provisions to achieve the physical clustering necessary for accommodating spatial data and for coping with fast access to spatial data [4-5]. DBMSs designed for commercial usage lack these provisions. Thus, they are not well suited for GIS.

The semantic DBMS that will be used in this GIS implementation does not contain these deficiencies. It has features that make it a high performance and efficient parallel semantic DBMS that is not only useful for most typical database applications, but is ideal for use in specialized domains such as earth science. This DBMS provides highly efficient accumulation and retrieval of general, scientific, and spatial data.

This semantic DBMS has been developed with a semantic/object-oriented approach. Thus, this system satisfies the three essential needs of many database applications: strong semantics embedded in the database to handle the complexity of the information, storage of multi-dimensional spatial, image, scientific, and other non-conventional data, and very high performance that allows rapid flow of massive amounts of data. Moreover, the semantic parallel architecture of this database system also provides efficient and flexible access to a large collection of data stored on various physical devices. Finally, data reference transparency is an inherent property of the semantic binary model system [2, 6].

3. Approach

This system was implemented in the UNIX environment and uses the X-Windows system to manage the display and the windows' management. The implementation language used was C++. The main software tool used for the graphics was Simple Raster Graphics Package (SRGP).

3.1 Hardware Description

The main computer used to compile, test, and run the GIS system is a SPARC Server-10 machine. This is a Sun workstation with 96 MB of RAM. In addition, this machine has 2 RISC based processors with 50 MHz each that give enough power to the computer to comfortably run the application.

4. Semantic Database

A semantic binary database schema is a set of categories, relations, and database types. A category is a specification for database abstract objects that belong to this category. Each category may have several relations with other categories and data types. A relation from a category to a data type is called an attribute, and a relation from a category to a category is called an abstract relation [2].

4.1 Schema Design

The spatial database used for the GIS has some specific requirements. This database has to store the spatial and semantic data together in the same database. Some GIS systems store the spatial data separately from the semantic data [7] which makes the system inefficient and difficult to use [6]. This database also has to store several types of information for many different data sets. For example, for each data set the name and description of the satellite, the instrument used, a color table (containing color, value, and a short description), the observation program used, the collection of all binary data files containing the data maps, and the date for each map in the database must also be stored.

Based on these requirements, the following semantic database schema was designed and implemented:

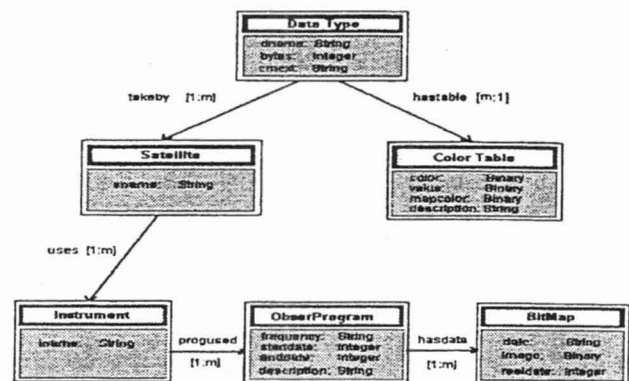


Figure 1. Semantic schema for the GIS's database.

5. Semantic GIS Implementation

The semantic GIS is completed independent from the data upon which it operates. It retrieves all the necessary information (data) from the semantic database, making the GIS a generic system that works for a variety of data sets.

5.1 User Interface

We endeavored to develop a friendly graphical user interface that is simple to use but still maintains a high level of flexibility and advanced data manipulation techniques. This system is windows based and interacts with users through the use of the keyboard and mouse. Hence, users have the flexibility of selecting any menu option using their preferred device. The main features of the GIS's interface are the top bar, drop-down menus, buttons, and icons. Through the use of these a user can easily manipulate the data, navigate through the semantic database schema, and retrieve the desired information. The menu options include projection, data, satellite, instrument, frequency, rotation, date, zoom, cloud, and exit. Their functions are described below.

- Projection - allows users to choose the desired type of view. The current available views include orthographic, orthogonal, sinusoidal, stereographic, homolographic, and mercator. Fig. 2 depicts a projec-

tion sequence of ozone data maps in the above mentioned sequence.




Figure 2. Ozone data maps in different projections.

- **Data** - allows users to choose the desired data set. This option performs a query to the database requesting the names of all the currently stored data sets. The result of the query is presented to the users as a list of possible options.
- **Satellite** - allows users to choose the desired satellite. This option performs a query to the database requesting a list of all recorded satellite names for the selected data set. The result of the query is presented to the users as a list of possible options.
- **Instrument** - allows users to choose the desired instruments. This option performs a query to the database requesting the names of all the instruments which were used in the satellite of interest. The result of the query is presented to the users as a list of possible options.
- **Frequency** - allows users to view the available frequencies for satellites and instruments. This option performs a query to the database requesting all recorded frequencies for the satellite and instrument of interest. The result of the query is presented to the users as a list of available frequencies.
- **Rotation** - allows users to 'fly over' a rotating earth and view changes in data values from longitude to longitude. There are three options available for this menu item. The increase option increases the rotation speed of the image by a factor of two. The reduce option gradually reduces the rotation speed by a factor of two. The stop option causes the rotation to cease completely.
- **Date** - allows users to view an animation of the data or a single frame at a time for a particular date. Every time a new frame is displayed, the corresponding measured date is also displayed on the lower portion of the screen. Among the options available for this menu item are single frame forward, single frame backward, reverse, increase, and stop.
- **Zoom** - allows the users to zoom in or out from a particular geographical location. For each selected area, a new window is opened displaying an animation of the data at the desired magnification rate. Magnification can range from 20% (zoom out) to 120% (zoom in). The main program window interface and all of the windows containing the selected zoomed frames

can be displayed on the screen at the same time. This greatly increases a user's observation and analysis capability since different geographical locations can be viewed on several windows at the same time at different magnifications and projections. Fig. 3 displays a sample of a screen after several areas have been magnified from an ocean temperature data frame.




Figure 3. Frames showing magnified areas of ocean temperatures data.

- **Cloud** - allows users to choose between two types of views: visible and non-visible clouds. The results of each of these views varies depending upon the type of data being viewed.
 - **Visible:** With ocean temperature data, this view allows users to see the cloud coverage that was recorded on the day that the observation was made. With ozone data, this view allows users to see the zones where no data was recorded due to instrument malfunction. With SeaWiFS data this option is not valid.
 - **Non-Visible:** With ocean temperature data, the frames are displayed without any visible clouds. This is done through the use of substitution. Basically, pixels representing clouds are substituted by previously recorded values. With ozone data, substitution is also used. Here, the last ozone reading is displayed in the areas that were cloud covered.

A feature of the semantic GIS's interface that is very rare in other GIS systems is the color table. Basically, each data set has a different color chart to aid in the interpretation of the colors. Once the user selects one particular data set, a color table (color, value) corresponding to that data set is retrieved from the database and displayed on the screen next to the map frame. For ocean temperature data, values are given in degrees Celsius. For ozone, values are given in Dobson Units (DU). Dobson units indicate the thickness of the ozone layer if it were measured at 0 degrees Celsius and at standard atmospheric pressure (1 atm) [8, 9].

Most of the commercial GIS and image processing packages on the market today do allow users to open data/image files, zoom, classify the data, and apply many other image processing techniques. They do not, however, provide users with specific information about the data such as the date read, the instruments and satellites used to collect the data, and many other textual and technical details about a particular data image. In addition, these packages do not allow users to specify a time frame (start-end date) in which they would like to view an animation of a particular data set. The semantic GIS, unlike most of the commercial packages, allows users to view how data values change over days, weeks, months, or years.

5.2 Display Process

The semantic GIS displays the spatial data image in static and dynamic modes. This section will concentrate on the static displaying process. The next section (data animation) will discuss the dynamic displaying process and analysis.

During the static display process, a data file for one particular date is retrieved from the database and placed in a buffer. This buffer is passed to a method where the whole data file is processed differentially depending upon the data type and projection. For instance, if the active projection is homolographic, the following mathematical computation needs to be performed before starting to process the values from the buffer:

```
xi = XGetImage(d_dpy, d_act.drawable.xid, hpw, 220, Dw, Dh,
             AllPlanes, ZPixmap);
for(int y = 0; y < Ph; y++)
    for(int x = 0; x < Pw; x++)
    {
        cx = R * (1 - ssin[y] * ccos[x] / (1 + ssin[y] * ssin[x]));
        cy = R * (1 - ccos[y] / (1 + ssin[y] * ssin[x]));
        if((f * (cx + cy * Dw)) < (Dw * Dh))
            map[x + Pw * y] = &xi->data[f * (cx + cy * Dw)];
    }
```

Once the computations for the projections are done, the process in the file buffer begins running a method called 'Next_Frame'. It is in this method that the image is computed. In order to produce the right image, the program needs to "remember" the location of cloud cover or instrument failure. Further, when ozone data is being display, the program needs to know the locations and outlines of the continents. The frame containing the outlines of the continents is superimposed over the ozone data frame resulting in two different frames being displayed at the same time. Hence, when the frame is finally displayed on the screen, users view the Ozone data and the continents outline all in one frame. This overlaying process is done to give users a better geographical orientation of the ozone layer values.

When all the required calculations and checks are done, the image is displayed. The algorithm executed in the 'Next_Frame' method does not run in linear time. Instead it requires a large portion of the total time it takes

to display the frames on the screen.

5.3 Data Animation

During the data animation process, more computation needs to be performed in order to achieve a dynamic view of all the data sets. If the date option (increase) and the rotation option (increase) are active, additional mathematical calculations need to be performed in the 'Next_Frame' method. These calculations would be used to rotate the whole image at the same time that the new frame is processed and displayed. Two time functions need to be executed during the flight and used to produce the combined result. The resulting output is a sequence of several rotating frames portraying data from nine or ten years (depending on the data set).

Our UNIX base GIS has a limited movie speed primarily due to the use of the X-Windows system. However, the database retrieval time (which is very small) and the computation time required to process the data before sending it to the screen (exponential) also plays a role in the speed of the animation. Therefore, the resulting animated movie has a speed of about five frames per second.

5.4 Data Retrieval

The semantic GIS is constantly retrieving information from the semantic database. For example, when the program is running in the dynamic mode (i.e., an animated movie), a request is made to the database every few milliseconds to retrieve the date and binary data values. The semantic database can easily cope with this high demand and still maintain a response time in milliseconds. However, how and when is the data retrieved from the database?

Each time that a menu option is activated, a query is performed to the database requesting a particular set of information. For example, when the satellite option is activated the following source code is executed to retrieve the data from the database and display it on the screen:

```
Var S = RangeQuery(TheRelation("DATATYPE::dname", DB,
                               dtname));
SetQuery Satel = S(TheRelation("DATATYPE::takeby", DB));
while( Satel.GetVarInc(sateln))
{
    Var SName = sateln["Satellite::sname"];
    sprintf(satel_name[i++], "%s", pChar(SName));
    SRGP_text(SRGP_defPoint(xl + 10, yl), satel_name[i-1]);
    yl += 20;
}
```

As can be seen from this source code listing, retrieving data to get a list of possible menu options is very simple. However, when a different data set is selected, and the user has already chosen all the options for a particular satellite, instrument, and frequency that he or she desires to view, more complicated queries are done to the database to extract the correct information and start the displaying process.

at different magnification levels and at different projections for as long as they desire. Although examples such as that shown in fig. 3 can aid in understanding and appreciating this system, one truly gains a full appreciation of it when working with it and applying GIS data manipulation techniques that are not easily described textually.

- Another feature of this system is its high level of flexibility. It is completed independent of the database upon which it operates. The system retrieves all necessary information (including the data format) from the database. Thus, different data sets can be displayed without having to change the system itself. The semantic GIS's flexibility is further illustrated in that users can click on a pull-down menu option, button, or icon at any time and a variety of operations will be performed instantly. One can view a different data set, view data on a particular date, zoom in over Europe, make a query to the database, or just stop the whole system until a decision is made to move on again.
- A high level of observation capability is an additional feature of this semantic GIS. Users can rotate a data frame in any selected projection until the desired view of the geographical location under research is achieved. Moreover, clients can use the zoom option to zoom in on several geographical locations and display them in different windows at different magnification levels along with the main user interface.
- Every displayed image is fully documented in this system. A description containing the data type, name, frequency, instrument, and satellite is given on the top border of each image. Additional documentation includes a color table giving a one-to-one relationship between colors and the values of the displayed data, a description of the color table, units values, and the date for each frame.
- The inclusion of an easy to use interface allows clients to navigate through the semantic database. Users do not need to have knowledge about the database schema or the information stored on the database. By simply choosing the available menu options, users are guided through all the categories and relations in the database. Further, they can retrieve all the desired semantic and spatial data at any time during the program execution.

The semantic GIS could also be considered an intelligent system. It is able to "remember" distinctive geographical information such as continents and clouds and then display this information in other projections as a point of reference for data precession or omit it for data clarity. In addition, this system is completely independent of the data upon which it operates and, since this GIS works with a semantic DBMS, the system is efficient and reliable.

Finally, another nice feature of this application is the animation/movie capability. Users can not only open a data file and view it, they can also view an animation/movie of a selected data set. Depending on the fre-

quency of the data under study, users can view how data changed over a period time through the animation, yet still stop the animation at any time to view the details.

7.2 Disadvantages

The main disadvantage of this system involves the display area. The animated movie displayed by the system does not achieve a very good quality movie speed. This is mainly due to the use of the X-Windows system to manage the display and windows. A minor disadvantage of this system is that it is not very portable.

7.3 Performance

The overall performance of this system is very good. The retrieval and transfer, time is almost instantaneous for textual information. The retrieval, transfer and display time for approximately 64KB of spatial data is 8-10 milliseconds. This results in an animation speed of about five to six frames per second. This system was tested with both the ATM switch and the regular Ethernet. As expected, the Ethernet took a few more milliseconds. Additionally, multiple users can efficiently run this application from several client computers and access a single centralized database.

7.4 Data Analysis

This GIS system has also allowed us to analyze, interpret, and study all the spatial data currently stored in the database. For example, ocean temperature data has been analyzed and the regions of high temperature water have been observed. One observation region close to the Tropic of Cancer in the Atlantic Ocean has warm water during the entire summer (June-October) making conditions favorable for the development of many hurricanes and tropical storms.

Ozone data has also been interpreted, analyzed, and observed. The development of the ozone hole (regions where the total ozone column is less than 200 DU) has been monitored. For example, figs. 4-6 depict a graphical sequence of the Antarctic ozone hole formation taken from monthly readings from the Nimbus-7 Satellite. In addition, fig. 7 portrays daily ozone data readings from the Meteor-3 Satellite on April 16, 1993, and fig. 8 displays ozone data from October 6, 1993. Fig. 7 shows the ozone layers at the beginning of the year and, as can be seen, all the values are greater than 230 DU. Fig. 8, on the other hand, depicts Dodson Unit values less than 100 DU over the South Pole (Antarctic). These values were recorded by NASA's TOMS flying on Russia's Meteor-3 satellite in October 1993. It was during this time that the ozone hole was 9 million square miles, slightly smaller than 1992's record ozone hole of 9.4 million square miles. On this day, a balloon sensor flown by NOAA measured 90 DU while a spectrophotometer on the surface measured 88 DU [12]. These two maps, from the same year and satellite but from different seasons, show how variable ozone layer values can be from one month to the next.




Figure 4. Ozone hole starts forming.




Figure 7. Ozone data at the beginning of 1993.




Figure 5. Ozone hole covers larger area.




Figure 8. Day of the largest ozone hole.




Figure 6. Ozone hole starts to diminish.

7.5 Ways to Improve the System

There are several procedures that can be performed to improve the GIS and achieve better system performance. Different tools could be used to handle the data display and windows management. Tools such as SRGP are gen-

erally good but are a bit primitive for our purposes. There are better tools and development kits on the market today that can be used to provide users with an easy to use interface yet still provide good graphic displaying techniques. This will render the program maintenance easy and greatly improve the animation.

In addition, more efficient algorithms should be developed to replace the existing ones that run in exponential time. The computation time will be greatly reduced if linear time algorithms are developed. These minor improvements will have a small but visible impact in the system's performance.

8. Conclusion

The design and implementation of the semantic GIS covered several areas. First, its semantic database was designed and created, and all the textual and spatial data was loaded into the database. Then the main areas of the system — the user interface, displaying process, data animation, data retrieval, and data transfer — were implemented with the help of some tools and devices. Finally, all

these components were tightly integrated to form a novel, innovative, and practical GIS.

The resulting system has facilitated the interpretation, storage, manipulation, analysis, and display of spatial data such as ocean temperature, ozone (TOMS) and simulated SeaWiFS data. At the same time, this system has helped in the testing process of the efficient, high performance parallel semantic DBMS being developed at the HPDRC.

As was mentioned in Section 7 (data analysis), with the help of the GIS system, the ocean temperature data has been analyzed and the ozone layer has been tracked for a period of more than 14 years. When is the ocean temperature near the west coast of Africa over 30 degrees Celsius (90 degrees F)? Did this occur when the ozone hole was first detected by the TOMS instrument? During which weather seasons is the ozone layer thinner and over what regions? These are only a few of the questions that can be answered using the semantic GIS system as the main tool for data examination, interpretation, and manipulation. This GIS guides users in retrieving all the desired information from the database and helps them to better analyze and observe the spatial data using a wide variety of menu options.

The semantic DBMS has also been tested. Database retrieval time has been monitored and retrieval and transfer time has been checked while running the GIS system using an ATM switch and Ethernet. Finally, data stored in the database has been periodically checked to ensure that the DBMS has been maintaining protection, security, and consistency. In sum, test results have shown that the semantic DBMS delivers efficient spatial and textual data retrieval time while maintaining consistency of the data.

8.1 Future Work

Future work includes the development of a Java version of the semantic GIS system. This task is currently underway and is, in fact, almost completed. The main purpose for implementing this change is to make the semantic GIS portable across several platforms.

The Java version of this system will be a client/server application that will be available through the World Wide Web. The client program will be a Java Applet consisting of a user interface, data display, and data manipulation. The server application will handle the data retrieval (database queries) and data transfer. All these components will be tightly integrated with the semantic database to achieve the system's goals and provide better features than the current system.

Acknowledgements

This research was supported in part by NASA (under grants NAGW-4080, NAG5-5095, and NRA-97-MTPE-05), NSF (CDA-9711582, IRI-9409661, and HRD-9707076), ARO (DAAH04-96-1-0049 and DAAH04-96-1-0278), DoI (CA-5280-4-9044), NATO (HTECH.LG 931449), and State of Florida.

References

- [1] N. Rische, W. Sun, D. Barton, Y. Deng, C. Orji, M. Alexopoulos, L. Loureiro, C. Ordonez, M. Sanchez, & A. Shaposhnikov, Florida International University High Performance Database Research Center, *SIGMOD Record*, 24(3), 1995, 71-76.
- [2] N. Rische, *A database design: The semantic modelling approach* (McGraw-Hill, 1992), 528.
- [3] M.J. Jackson, W.J. Jame, & A. Steven, The design of Environmental Geographic Information System, *Philosophical Trans. of the Royal Society of London*, 324, 1988, 77-84.
- [4] A.V. Frank, Requirements for a database management system for a GIS, *Photogrammetric Eng. and Remote Sensing*, 54(11), 1988, 1557-1564.
- [5] T. Keating, W. Phillips, & K. Ingram, An integrated topologic database design for Geographic Information Systems, *Photogrammetric Eng. and Remote Sensing*, 53(10), 1987, 1399-1402.
- [6] N. Rische & Q. Li, Storage of spatial data in semantic database, *Comp. in Eng. 1994*, Proc. of the 1994 ASME Int. Computers in Eng. Conf., Minneapolis, MN, September 11-14, 1994, ASME, 793-800.
- [7] T.C. Waugh & R.G. Healey, The GEOVIEW design: A relational database approach to geographical data handling, *Int. J. of Geographical Info. Sys.*, 1(2), 1987, 101-118.
- [8] J.R. Asker, US/Russian satellite maps: Worst Antarctic ozone hole, *Aviation Week and Space Tech.*, 139(17), 1993.
- [9] http://rtfm.mit.edu/pub/usenet/news.answers/Ozone_depletion/intro
- [10] C. Webster, Desegregated GIS architecture: Lessons from recent developments in multi-site database management systems, *Int. J. of Geographical Info. Sys.*, 2(1), 1988.
- [11] G. Muffin, Raster versus vector data encoding and handling: A commentary, *Photogrammetric Engineering and Remote Sensing*, 53(10), 1987, 1397-1398.
- [12] W. Gregg et. al. SeaWiFS technical report series, *NASA Technical Memorandum 104566*, 15, January 1994.

Biographies

Elma Alvarez is a Senior Systems Analyst at the American Bankers Insurance Group in Miami, Florida. Previously, Dr. Alvarez worked at Florida International University and FLP. She is involved in research in the areas of OO database systems, geographic information systems and multimedia. Dr. Alvarez is a member of the IEEE and the IEEE Computer Society.

Naphtali Rische is a Professor of Computer Science at Florida International University. He earned his B.Sc. and M.Sc. degrees in computer science from the Israel Institute of Technology, and his Ph.D. degree in computer science from Tel Aviv University in 1984. Dr. Rische's current research focuses on efficiency and flexibility of databases, distributed DBMS, high-performance systems, database design tools, and internet access to databases. He is the founder and director of the High Performance Database Research Center at FIU, and is a member of the New York and Florida Academies of Sciences, ACM and IEEE.

David Barton is a Professor of Computer Science at Florida International University. He earned his B.Sc. degree from the University of London and his Ph.D. degree from Cambridge University in 1966. His research interests

include distributed systems and data communications. Dr. Barton has published more than twenty papers on various topics in computer science. He is a member of the ACM, and is a Fellow of the Royal Astronomical Society and the Cambridge Philosophical Society.

