Real-Time Database System for Meteorological, Hydrological and Oceanographic Data

Esa Falkenroth
MHO Data Warehouse,
Swedish Meteorological and Hydrological Institute
SE-601 76 Norrköping, Sweden

Abstract. The Swedish Meteorological and Hydrological Institute, SMHI, has developed a new real-time database system that combines meteorological, hydrological, and oceanographic data in a single framework. The system is based on a commercial object-relational database system (Informix Dynamic Server/UDO) and a distributed object computing infrastructure (CORBA/C++/Java). This paper presents the current system with a focus on issues related to the system architecture and the unified data model for weather and water information. As a conclusion, the experience of building a large-scale object-oriented database system is presented along with some performance measurements.

1. INTRODUCTION

The Swedish Meteorological and Hydrological Institute provides information for businesses and activities that depend on weather or water. The core of this information is meteorological, hydrological, and ocenographic observations, forecast fields, images, *etc.* (also known as MHO data). To assemble customer products based on MHO-data in a more flexible and economical manner, SMHI has initiated an extensive effort towards rationalisation of the production processes [Lar99]. A central part of this effort is a new infrastructure with an object-relational database system, ROAD. This paper presents the ROAD database system, its data model, and the interfaces to the data acquisition systems and the production process.

Production of MHO data

At SMHI, the production of MHO data is based on a process-oriented information philosophy. Hence, MHO data are processed in three separate stages: co-production, end-production, and distribution.

The purpose of the co-production is to produce quality-controlled MHO-data in real-time. This includes data capture from observation stations, satellites, radars, ships, and airplanes. A wide range of numerical models are also included in the MHO data. To achieve a single perspective on the weather situation, observations and forecasts from several sources are combined in a coherent data set that describes current and future weather and water situation as accurately as possible. Today, the coherent model will be assembled by a meteorologist through an forecast control process [Lar99] where different parameters in observations and forecasts are combined by weighted interpolation. A typically coherent model will consist of data from MESAN (mesoscale analysis) [Häg97] for the current weather, various HIRLAM [Käl96][Gol95] forecasts for the short-term weather, and medium-range forecasts from ECMWF. Known deficiencies of forecast models can be manually compensated through graphical editing tools. In a sense, the editing and interpolation of models resembles a meteorologist's effort to build a mental model of the weather.

The result of the co-production is a single accredited high-quality representation of observations and forecasts in the real-time database.

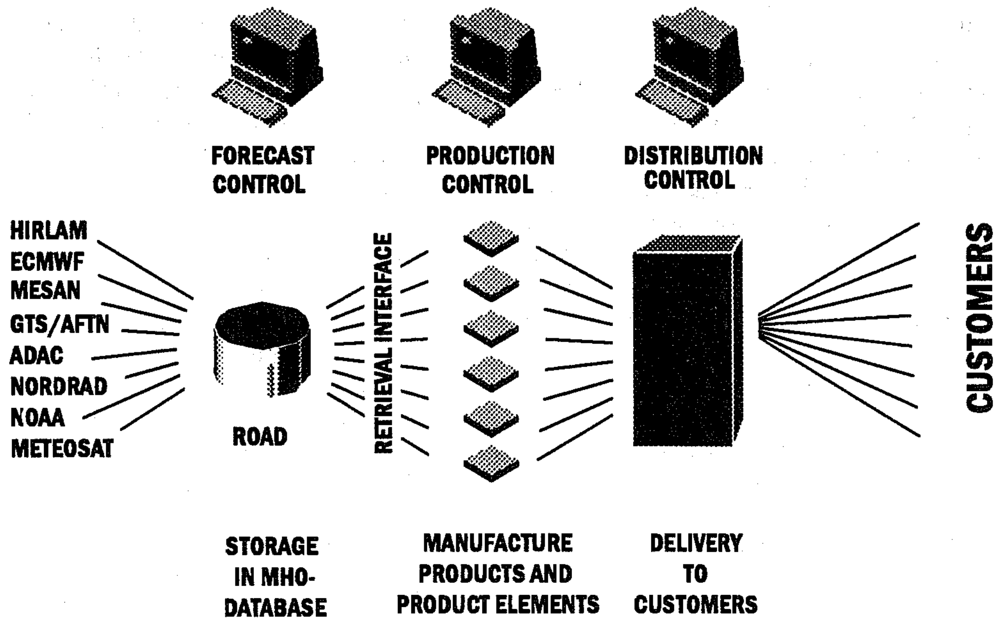

Figure 1.

The next stage is the end-production, which creates descriptions of major weather and water phenomena related to geographical areas and time intervals. Here, the co-production data forms the basis for a semiautomatic processing and assembly of observations and forecasts into MHO-products according to customer demands.

2. Database-centred architecture

In the production system for weather and water information, SMHI have adopted a database-centred approach: all MHO data is collected in an Informix Dynamic Server [Inf97]. The result is a clean, simple, and consistent interface with provides the necessary flexibility, consistency, parallel data access, data integrity, and authorisation mechanisms.

The production system involves five different stages:

- Data capture with quality control
- Storage of real-time data and archival data
- Tools for manual and automatic production
- System for distribution of products to customers

These are illustrated in Figure 1. A central component of an efficient and flexible end-production is the real-time database system. All captured data is stored in the database to facilitate retrieval and processing.

3. Conceptual data model

The ROAD data model combines meteorological, hydrological, and oceanographic data in a single framework. The rationale for unification is the reduction of redundant data storage, and the introduction of a standardised cross-enterprise interface for data access. In the long run, we also hope that a unified data system will promote synergies between the research disciplines of meteorology, hydrology, and oceanography. .

The conceptual model is generic, which means that the model is not limited to storage formats such as SYNOP, BUFR, or GRIB [STA89]. In contrast, each value is an individual entity and the model allows new types of data and storage formats to be added. The goal of this data model is to provide a flexible structure for organising MHO data. By a flexible structure we mean a data model that provides appropriate abstractions and an independence from data sources, applications, and the organisation of SMHI.

Essentially, all MHO-data belong to the category spatio-temporal data. This means that all values are associated with time point (or time interval), and a geographical point (or area). Beside the spatial and temporal reference, each value is described by the following attributes: source, MHO-type, quality, and method.

- Each MHO datum is associated with a time interval that describes when the data is valid. The valid-time is described as an intervall between two UTC timepoints.

- Each MHO datum is associated with MHO-type-information: temperature, wind speed, wind direction, precipitation, *etc.* (more than 600 different MHO-types). In turn, the type identifies the unit.
- Each MHO datum is associated with spatial reference. The reference is either a single point, a line, a circle, or a polygon area. Polygon areas allow arbitrary irregular areas such as cities or counties to be represented. The spatial data is described with latitude-longitude coordinates. In addition to the geographic description, each MHO datum is associated with a level that states the height where the datum is valid. Levels can be expressed in many different ways: meter above sea level, meters above ground, level number in forecast fields, standard pressure levels, etc.
- Each MHO datum comes from a source. The source can be a station (automatic station, airport, radar, satellite, ship), or a computation (numerical model).
- Each MHO datum is associated with a quality description.
- For calculated values, the algorithm (method) is identified.
- Finally, the actual value is stored as a numeric value (IEEE-float), a text, an image, or a two-dimensional matrix of IEEE-floats.

4. Data in ROAD

To be effective and useful for the full range of users, the ROAD database must contain real-time data with suitable temporal and spatial coverage, adequate accuracy and resolution, and relevant parameters. Therefore ROAD continuously receives MHO data from satellites, radar, observation stations, and computing centres for weather forecasts.

Image data

ROAD stores satellite images and radar images in their BLOBS (Binary Large Objects) using their original format (NORDRAD, TIFF and GIF87). Along with the standard meta-data (spatio-temporal reference, source, quality, method, and MHO-type), ROAD stores information about the image (resolution, area designation, and projection).

Currently ROAD stores image data from Prosat (Meteosat and NOAA), two swedish Nordrad radars, and forecast images from Metview.

Forecast Data

ROAD stores several types of forecast fields and postprocessed fields: ECMWF-GLOBAL, ECMWF-EURO, HIRLAM22, HIRLAM44, HIRLAM-1D[GOL95], and MESAN. The contents of GRIB-fields [Sta89] are unpacked and decoded and stored as IEEE-floats. Currently, the nordic area is represented along with adjacent drainage basin; however, some global wave data is also stored.

Observational data

ROAD also stores observations from manual and automatic stations, airports, ships, aeroplanes and aerological probes. Most of this data is delivered through GTS and the national network for automatic and manual stations (ADAC). GTS provides METAR, TEMP, PILOT, CLIMATE, AIREP, SHIP, CYCLONWARNING, ICEBERGFORECAST, SIGMET, QBC, and TAF.

A coherent forecast (PMP)

Besides raw data and forecast fields, road also stores interpretations on the weather situation, which are used in the automatic production. Based on the available forecasts, a new postprocessed forecast is created. With a 0.2 degrees grid resolution the nordic area is tiled into a 102 × 116 matrix. For each grid element, 35 postprocessed parameters are stored. These are called product montage parameters (PMP). An example of such a parameter is weather symbols and forecast text for specific areas. These parameters provide a fundament for assembling forecast products, warnings, and weather reports for internet and WAP.

Web data

ROAD also stores web data using a Web datablade [Inf99]. This is mainly used for storing metadata: description of data sources, parameters, geographic areas, quality, *etc.*

5. Access to data

The MHO data is accessed by applications through SQL. Currently, interactive users, dynamic SQL, JDBC, and ODBC are only allowed in special cases since they may jeopardise the performance and stability of the database system.

The interaction with the database system are summed up by the following two categories: requests to store observations or fields, and retrieval of data initiated by applications. These categories are represented by two sets of APIs (application program interfaces): loader APIs and query APIs.

Application program interfaces for data loading

Data capture is performed by several different servers. Through CORBA (Objectbroker and Orbix) and other TCP/IP based solutions, these servers deliver data to a set of loader processes that execute on the same server as the database management system. In turn, these load processes insert data into tables in the database.

Application program interfaces for retrieval

Data can be retrieved at several levels. Either an application retrieves an observation or a complete field. A more advanced retrieval requests an automatically generated text that described the weather at a certain position and time. Regardless of the request, the application developers do not need to know how the data is structured and stored in the database. Instead, they identify the desired parameter, position, and time range.

In the long run, the retrieval layer will allow developers to leverage from each other's effort.

6. META DATA

At the lowest level information is organised using parameters that represent individual observations or fields. To facilitate access, the data are also organised in information-bases (IB). An IB represents information about data sources, life cycle time for data, geographic areas, *etc*. The IBs represent the table of contents for the database.

7. IMPLEMENTATION

The ROAD database is based on a commercial object-relational DBMS, *Informix Dynamic Server* with Universal Data Option 9.14UC6X7 (soon to be upgraded to *Informix Dynamic Server 2000* (9.20). It is extended with Geodetic Datablade 2.21 for geospatial data [Geo97], and web data (Web Datablade 3.32). Currently, the system runs on a SUN Enterprise Server 3000 with six ultraSparc2 CPUS and 1.5 gigabyte main-memory. It runs on Solaris version 2.6. Dual disk arrays (Network Array A5000) provide 75 gigabyte of mirrored diskspace. An FDDI ring connects the database server with other servers for data sources and applications.

Incoming data from GTS/AFTN and reports from automatic stations are received by an Alpha/VMS where they are unpacked, buffered, and sent to load processes executing on the database server using Object Broker (ver. 3.0). The load processes store the information using ESQL through shared memory. Radar and satellite images are loaded through TCP/IP interfaces (socket/rcp). Forecast fields are made available through CORBA and NFS.

Captured raw data is stored in 7 tables in the database server:

- Numvalue
- Stringvalue
- Longstringvalue
- Forecastvalue
- Field
- Radarimage
- Satelliteimage
- Mobileship

In addition, data related to the coherent forecast is stored in 2 tables: Metviewimage and PMP area.

The data model is fully normalised and uses object-relational features to provide inheritance, and type extensibility. This means that observational data are not stored in telegram format. Instead, the telegrams are decomposed and each value is stored separately. The values are identified by ROAD-specific parameter number. This simplifies interfaces and application programs.

Connections to production applications are mediated through CORBA (ORBIX). An additional layer provides support for retrieval of abstract weather information. For instance, it is possible to state forecast queries that return automatically generated texts that describe the weather at a particular point for a particular time. To be compatible with older applications, a set of APIS also provides access to fields and observations in the original telegram and GRIB formats.

8. EXPERIENCE

Building a real-time database system for MHO-data is a complex undertaking. Here, an essential success factor is clear requirement specifications. Equally important is a "strategy for performance" that monitors how changes in the requirements affect the performance of the database system. The performance strategy should enforce the construction of scalable data structures and APIS, for example, use of thread safe APIS for database operations.

Some of the encountered problem were unanticipated. For instance, it was very hard to construct methods and algorithms that analysed forecast fields and generated abstract weather symbols, texts, and probabilities.

The choice of DBMS (Informix IDS/UDO) has provided a flexible platform both in terms of modelling capabilities (object-relational modelling, inheritance and polymorphism), and extensibility (DataBlades, and user-defined routines)[Dbl97].

Naturally, a project of this size will encounter problems both with hardware and software. The Sun hardware have had problems with frequently failing GBICS for A5000 disks and low availability for HA-cluster. The Informix software have had unexpected problems with indexes and smart BLOBS in IDS/UDO. Most of these problems have been solved by Informix and Sun during the project. Today, the DBMS is operational and provides excellent performance.

Advantages

Besides all the expected advantages of database technology: safe concurrent access to shared data through transaction processing, optimised access through query optimisation and index-methods, independence from data formats and data structures, persistent storage of data with backup facilities, protection against data corruption at software and hardware failures, standardised interfaces for data access, authorisation mechanisms, *etc.*, the database system provided additional advantages in terms of program development. The major advantage is the unified view on MHO data: one unified model for structuring MHO data, a unified set of parameters (same units and same interpretation of parameters regardless of source), and a unified retrieval-interface for applications. The unification provides a good fundament for combining data in new easy ways.

This approach lowers the development costs and shortens the lead time for new applications. By gathering common functionality in the database system and its loader-processes, the application programs can be made simpler: less redundant code for decoding formats, less redundant checks of data, less problems comparing data of similar type, etc. The higher level of abstraction provided by the retrieval interface also contributes to simplified application development.

The distributed system architecture has been successful: the data can be retrieved from different platforms using the same industry-standard communication protocol, CORBA.

9. Performance evaluation

The purpose of the database is to provide MHO-data of highest possible quality to the production systems. This requires continuous updates with current observations and forecast field. Every day 30 gigabyte data are loaded into ROAD. This includes 2,200,000 rows of observation data, 70,000 fields, 200 radar images, and 50 satellite images per day. The performance is acceptable. During normal load, 90 GRIB fields (approximately 23,000 grid points) are loaded in 2 minutes. The maximum measured I/O rate for field loading is 6.6 megabyte per second which translates into 250 fields per minute. At peak load, the fields are loaded in 5 minutes. For large deliveries of synoptic data, the situation is more critical. Today the 6Z and 18Z takes approximately forty minutes to load. An inspection of the indexes and the loading software revealed some inefficient solutions and after a number of improvements we now measured a loading capacity in the range of 300 observations stored per second.

The time to retrieve a complete GRIB field is approximately 0.01 seconds; however, the retrieval of a set of individual grid points is slower. Here, we have constructed a computational indexing method that will drastically reduce the access time [Fal96][Fal99]. A new field-API sports an improved algorithm for extracting gridpoint. Given the new algorithm, it takes 0.18 seconds to retrieve four adjacent grid points and less than 50 seconds to construct a time series from 8000 different fields (6 milliseconds per gridpoint). For observations, the time to retrieve an individual observation is in the same range: approximately 0.01 seconds.

Performance of the A5000 disk subsystem is very high. Less than 1 procent of the CPU time is spent waiting for disk I/O. Instead, the limiting factor is the CPU capacity. Currently, the memory consumption is 1.8 gigabyte (Solaris operating system, the DBMS and the associated loader processes).

10. STATUS FOR ROAD VERSION 1.0

ROAD version 1.0 demonstrates the feasibility of a database-centric approach to handling meteorological, hydrological, and oceanographic data. The common retrieval interface to data allows more efficient application development with shorter lead times. The database system is operational as of October 1999, and the general concept of a real-time MHO database is currently being introduced in a step-wise manner at SMHI.

During March 2000, one thousand automated products are expected to begin routine production. Additional products will be introduced progressively over the next year. At the same time older existing systems will be migrated to use the central MHO database system. Meanwhile, we expect that there will be many additions to the current ROAD database system.

At this time, the system is limited by the available CPU-capacity (6 UltraSparc CPUS); however, the system is still new and there are ample opportunities to improve performance. New more efficient data representation and new APIs will address some of the performance issues uncovered by ROAD version 1.0. Once the performance has been improved, we wish to improve the spatial and temporal resolution of the data. Over the next years the extent of loading (field, radar, and satellite) will be increased with an additional 200 gigabyte per day.

Later, the system will be extended with archival facilities that store a subset of the real-time data on tape (in a robot archive). The goal is to provide transparent access to data regardless of its age and location (ROAD or archival system).

11. Summary

This paper describes the new real-time database system, ROAD, deployed at SMHI, Sweden. The purpose of ROAD is to store meteorological, hydrological, and oceanographical data in a single database. The construction of the database system relies on several items:

- Data sources (observations, radar, satellites)
- A wide range of numeric models
- Data acquisition with built-in quality control procedures
- Integrated communication (CORBA, TCP/IP, NFS, *etc.*)
- Extensible functionality of query language
- Support for geospatial and temporal data
- Smart storage of binary large objects (BLOBS)
- Meteorological, hydrological, and oceanographic expertise

The current system is based on a commercial object-relational database system, Informix Universal Server. It runs on a Sun Enterprise Server 3000 which has a capacity of storing 30 gigabyte MHO-data per day.

Acknowledgements

The MHO database concept was initially developed by I. Hedlund, M. Källström and P. Emanuelsson. The ROAD system was developed by a team consisting of personel from SMHI/Is. Two consultants have developed significant parts of the system: P. Hagström and M. Hedberg. System administration has been ably provided by SMHI/Id.

References

[Dat97] *DataBlade API Programmers Manual*, Version 9.1, Part No. 000-4812, Informix Press, March 1997

[Fal96] Falkenroth, E., *Computational Indexes for Temporal Data*, in Proc. 8th Intl. Conf. On Statistical and Scientific Database Systems, Stockholm, Sweden 1996

[Fal99] Falkenroth, E., *Database Technology for Control and Simulation*, Ph D Thesis, Linköping University, Sweden (in print)

[Geo97] *Informix Geodetic DataBlade Module*, User's Guide, Version 2.1, Part No. 000-3711, Informix Press, January 1997

[Gol95] Golvik, S., Bringfelt, B., Perov, V., Holtslag, A.A.M., *Experiments with nonlocal vertical diffusion in HIRLAM*, Hirlam Technical Report No. 18., SMHI, 1995

[Häg97] Hägglund, L., Ivarsson, K.-I., Olofsson, P.-O., *MESAN, Mesoskalig analys*, SMHI RML-series Nr 75, SMHI, Sweden March 1997 (in Swedish)

[Inf97] *Informix Universal Server*, Version 9.01, Part No. 000-3678, Informix Press, January 1997

[Inf99] *Informix Web DataBlade Module*, Application Developer's Guide, Version 4.0, Part No. 000-5474, Informix Press, 1999

[Ion98] *Orbix 2*, Orbix Programmer's Guide, Iona Technologies PLC, October 1997

[Käl96] *Hirlam Documentation Manual*, System 2.5, Ed. Källén, E., SMHI, June 1996

[Käl98] Källström, M., *Road Principer, fakta och idéer*, PM 1998-06-17, SMHI, June 1998 (in Swedish)

[Lar99] Larsson, A., *Short status report on the project: Rationalisation in the Production Process – RiPP*, in Proc. 10th EGOWS meeting, De Bilt, Netherlands, 7-10 June 1999

[Lin99] Lindh, O., *Experiences from the new Forecast Production System in Sweden*, RIPP Version 1.0, in Proceedings of 4th European Conference on Applications of Meteorology, ECAM'99, (CD-version), Norrköping, Sweden September 1999.

[Sta89] Stackpole, J.D., GRIB and BUFR – *The only codes you will ever need*, in Proceedings of 2nd Workshop on Meteorological operational Systems, ECMWF, December 1989.