

Using the CALANDA¹ Time Series Management System

Werner Dreyer, Angelika Kotz Dittrich, Duri Schmidt
 {dreyer, dittrich, schmidt}@ubilab.ubs.ch
 UBILAB, Union Bank of Switzerland, Zurich

In financial institutions, time series are important for economic and financial research as well as for various non-research related activities like portfolio management. In general, scientific applications make extensive use of time series. Using traditional general-purpose DBMS proves inadequate for the requirements of this domain. We show how to cope with the problem that a DBMS for time series management must provide a complex data model and at the same time a user-friendly interface. We demonstrate how the mapping between the components of the data model and the elements of the interface is done in the time series management system CALANDA. CALANDA has been designed under the assumption that application specialists will use it rather than computer experts. Therefore, its data model is accessed via a graphical interface which users familiar with spreadsheets or graphics-based database software need very little education to use.

The main characteristics of the CALANDA system and its data model [DKS94] are:

- It is a special-purpose DBMS with full-fledged storage, retrieval and manipulation functionality for the time series domain.
- It has an object-oriented data model with root classes *Timeseries* and *Group*.
- Multivariate time series with query capabilities and time scale conversion are the basic abstraction.
- Groups are an effective categorization and aggregation instrument.

The interface to CALANDA has been designed according to three major principles:

- It is an interface with intuitive graphical elements and a menu-based way of manipulation. The elements of the data model are directly mapped to graphical components, the operations applicable in the model are represented by menu items.
- The interface has been designed to resemble well-known paradigms like spreadsheets or 4GL tools for relational databases.
- Tools in this interface serve two purposes: They provide an overview of the data base and they allow to retrieve and/or manipulate individual elements. Tools of these two kinds exist for meta data, that is, information on classes, as well as for primary data, namely time series and groups.

CALANDA has been implemented using object-oriented technology. As the underlying platform, we used ETOS, a seamless integration of the object-oriented DBMS ObjectStore and the C++-based application framework ET++ [WG 94]. ETOS provides persistent storage of objects, basic database functionality

¹ CALANDA is the name of a mountain (2805m) in the Grisons, Switzerland.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association of Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

SIGMOD '95, San Jose, CA USA

© 1995 ACM 0-89791-731-6/95/0005..\$3.50

as well as a framework for application and interface building. CALANDA runs under Unix on Sun workstations.

Our early experiences with economic researchers have shown that this system really meets their needs as to easy handling and straightforward but powerful manipulation of time series. Application examples are large company databases with historical security prices or macro economic databases containing raw and aggregated data on gross national product, employment rate and many other values.

The graphical user interface provides the user with a number of tools for activities like navigating through time series bases, browsing individual groups and time series, querying time series, defining new time series and group classes etc. Two examples of such graphical tools are shown in figures 1 and 2. Fig.1 presents a browser that shows the group hierarchy of a time series base and enables the user to navigate among groups, subgroups and time series. Fig.2 presents a browser and editor tool for time series showing the header attributes and a scrollable event list.

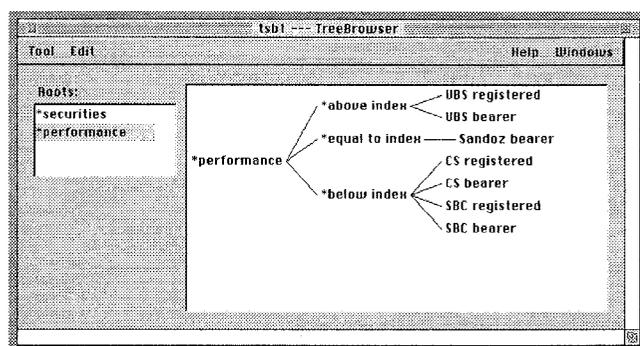


Figure 1: Browsing the group hierarchy of a time series base

timestamp	open	close	daily_vol
10.Aug.94	9811	7392	905
11.Aug.94	5038	7775	8732
12.Aug.94	8181	1370	9531
13.Aug.94	5144	9105	6742
14.Aug.94	4839	8036	7437
15.Aug.94	834	2659	8832

Figure 2: Browsing a time series

[DKS 94] W. Dreyer, A. Kotz Dittrich, D. Schmidt: An Object-Oriented Data Model for a Time Series Management System. Proceedings of the 7th International Working Conference on Scientific and Statistical Database Management (SSDBM'94), Charlottesville, Virginia, Sep. 1994.

[WG94] A.Weinand, E.Gamma: ET++ - a Portable, Homogeneous Class Library and Application Framework. In W.R.Bischofberger, H.-P. Frei (eds.): Computer Science Research at UBILAB, Proc. of the UBILAB Conference 1994, Zurich, Universitätsverlag Konstanz 1994.