The Multidimensional Database System RasDaMan

P. Baumann, A. Dehmel, P. Furtado, R. Ritsch, N. Widmann
FORWISS (Bavarian Research Center for Knowledge-Based Systems)
Orleansstr. 34
D-81667 Munich, GERMANY
Tel. +49-89-48095-200

{ baumann | dehmel | furtado | ritsch | widmann }@forwiss.tu-muenchen.de

1. ABSTRACT

RasDaMan is a universal – i.e., domain-independent – array DBMS for multidimensional arrays of arbitrary size and structure. A declarative, SQL-based array query language offers flexible retrieval and manipulation. Efficient server-based query evaluation is enabled by an intelligent optimizer and a streamlined storage architecture based on flexible array tiling and compression.

RasDaMan is being used in several international projects for the management of geo and healthcare data of various dimensionality.

1.1 Keywords

Array DBMS, multidimensional discrete data, raster data

2. INTRODUCTION

Arrays of arbitrary size and dimension, so-called *Multidimensional Discrete Data* (MDD), span a remarkably rich manifold of information – from 1-D time series and 2-D images to OLAP data cubes with maybe dozens of dimensions, with sizes from a few kilobytes to several Gigabytes, as spatio-temporally discretized natural phenomena or as artificially generated data sets. MDD appear in a host of database applications: OLAP, statistics, earth and space sciences, medical imagery, wind channels, simulations, and multimedia comprise but a few representatives.

Various research activities have been reported, however focused on particular application fields like healthcare imagery [2], geo data [11], or OLAP [1].

The RasDaMan array DBMS, the outcome of a European ESPRIT Long-Term Research project [12], has the goal of establishing arrays as first-class database citizens. As opposed to other R&D work, RasDaMan strives for domain-independent support for arrays of arbitrary size, dimension, and base type through a general-purpose declarative query language paired with internal execution, storage, and transfer optimization. This enables, for example, to maintain a data warehouse simultaneously with satellite images in a geo clearinghouse. A further difference to most other scientific work in the field is that the RasDaMan DBMS is fully implemented and in practical use in several projects.

In this contribution, we present the RasDaMan conceptual model and system architecture and two sample projects where the system is being used.

3. CONCEPTUAL MODEL

The RasDaMan conceptual model centers around the notion of a multidimensional array of arbitrary dimension, extent in each dimension – whereby each lower and upper bound can be fixed or variable –, and base type. Usually such an array will be an attribute of some other object, e.g., the "raw data" accompanied by "registration data" within an image. The operation set is based on RasDaMan Array Algebra [5] which allows for declarative expression of operations up to the complexity of the Discrete Fourier Transform. Array expressions are embedded into standard SQL-92 in the array query language *RasQL*. Operations include

- trimming (rectangular cutout) and section (extraction of a lower-dimensional hyperplane);
- induced operations which apply cell operations simultaneously to the whole array;
- generalized array aggregation;
- format converters to accept and deliver arrays not just in the client's main memory format, but additionally in a number of standard data exchange formats.

4. SYSTEM ARCHITECTURE

The RasDaMan client/server DBMS completely has been designed in an object-oriented manner, using standards wherever possible. The RasDaMan API consists of array-extended SQL-92, *RasQL*, and an ODMG [7] conformant C++ API, *RasLib*.

RasDaMan follows the classical two-tier client/server architecture with query processing done completely in the server. Internally and invisible to the application, arrays are decomposed into tiles, i.e., rectangular parts, which form the unit of storage and access.

The RasDaMan server is designed as a middleware which maps the array semantics to a simple 'set of blob' semantics where each tile is represented by a blob. The underlying DBMS, currently O2 [10], is operated as a persistent store which essentially only must offer a blob concept. Due to the easy portability of RasDaMan between different base DBMSs and storage systems, it easily integrates into existing IT infrastructures.

The query evaluator parses a *RasQL* query and builds an operator-based query tree. This query tree is optimized in two steps, algebraic query rewriting and second physical optimization based on tiling, clustering, and device information [6]. Part of this is to set up a sequence of operations on tiles and to rearrange tile access sequence to determine the cost-optimal evaluation sequence. A multidimensional index helps to identify the tiles involved in a query and to calculate the costs to retrieve them. Further internal components manage the schema catalog and the storage device characteristics used for cost-based query execution planning. Among the utility set available are *RasDaView*, a visual frontend for n-D array retrieval and manipulation (see Figure) and the *RasDL* schema processor.

5. APPLICATIONS

5.1 Human Brain Database

The European Computerized Human Brain Database (ECHBD) is a web database of structural and functional information about the human brain [8]. It contains 3-D images (modalities) from different sources and in different resolutions. Among the modalities maintained are regional cerebral blood flow with a resolution ±2.2 mm, anatomical MRI with voxel size 1x1x1 mm, and receptor-architectonical data with a resolution of 15 to 50 µm. Images from several hundred brains, both living and postmortem, have already been acquired. Each image occupies about 7.8 MB uncompressed. Based on a reference brain, images are spatially normalized during insertion. Once the brain cubes are in canonical form, areas can be addressed through voxel coordinates uniformly across all brains and modalities. Stored in a Web-accessible RasDaMan database, researchers can then run a variety of analyses, such as data comparison and similarity search.

5.2 Continental Drilling Program

The International Continental Scientific Drilling Program (ICDP) supports several earth science research projects where drilling is an essential [9]. The program aims at helping scientists from different countries to cooperate and link together their scientific data by providing a uniform infrastructure for scientists to register data directly at the bore hole. In an on-site data center, they can digitize and annotate drill cores piecemeal as they come in from the hole. With a diameter of 10 cm, a depth between 3,500 m and 5,000 m and a scan resolution of 103 dpi, a typical complete image of a drill hole occupies between 17 and 24 GB uncompressed. In parallel with such a 2-D bore hole image, 1-D sensor data can be measured, such as magnetism. Each drilling site, among them Hawaii, Mexico, Siberia, Japan, and Germany, maintains its own web accessible RasDaMan database. Additionally, all data will be phased into a central RasDaMan database.

Typical queries will include access to a particular depth, search for annotations, and search for particular geophysical phenomena.

6. CURRENT STATUS AND FUTURE

The RasDaMan system as demonstrated at SIGMOD'98 supports arrays of arbitrary dimensions over primitive and user-defined cell types. The client/server system is operational on Unix and Windows NT. The C++ API *RasLib* comprises both schema and instance operations. *RasQL* (see Figure) offers trimming, section, induction, format conversion, and the boolean aggregates `some_cells` and `all_cells`. The *RasDL* processor creates dictionary entries and C++ header files out of schema definitions. More than 60 query optimization rules are arranged in five selectable levels. Tile-based storage management encompasses configurable tiling and transparent tile compression.

The RasDaMan DBMS is being marketed commercially, in particular for the aforementioned geo and healthcare areas.

Next development steps encompass extending the operational power with dimension hierarchies and generalized aggregation. Further, the issue of internal compression in conjunction with efficient query evaluation will be investigated deeper. A storage layout language is under development to exactly control tile management and other tuning parameters. Last but not least, study of application requirements assisted by system benchmarks will be continued.

7. ACKNOWLEDGMENTS

The RasDaMan project is sponsored by the European Commission under grant no. 20073. We gratefully acknowledge the input of our end users in the RasDaMan project, Hospital General de Manresa (Spain) and the Spanish National Geographic Institute.

8. REFERENCES

[1] R. Agrawal et al. Modeling Multidimensional Databases. In Proc. ICDE'97 (Birmingham, April 1997), 232-243.

[2] M. Arya, W. Cody, C. Faloutsos, J. Richardson, and A. W. Toga. A 3D medical image database management system. Com. Med. Image & Graph. 1995, 210-214.

[3] F. Bancilhon and C. Delobel and P. Kanellakis. Building an Object-Oriented Database System. Morgan Kaufmann, San Mateo-California, 1992.

[4] Baumann, P. et al. Geo/Environmental and Medical Data Management in the RasDaMan System. In Proc. VLDB'97 (Athens, August 1997), 548-552.

[5] Baumann, P. The RasDaMan Array Algebra. RasDaMan Technical Report for012, FORWISS, 1998.

[6] Baumann, P., and Ritsch, R. Optimization and Evaluation of Array Queries in RasDaMan. RasDaMan Technical Report for013, FORWISS, 1998.

[7] R. Cattell. The Object Database Standard: ODMG-93. Morgan Kaufmann, San Mateo-California, 1996.

[8] European Computerized Human Brain Database homepage. http://www.atlas.neuro.ki.se/Hba/

[9] International Continental Scientific Drilling Program homepage. http://www.gfz-potsdam.de/pb5/icdp/

[10] O2 Tech., France homepage. http://www.o2tech.fr/

[11] J. Patel et al. Building a Scaleable Geo-Spatial DBMS: Technology, Implementation, and Evaluation. In Proc. ACM SIGMOD'97 (Tucson, May 1997), 336-347.

[12] RasDaMan homepage. http://www.forwiss.tu-muenchen.de/~rasdaman/

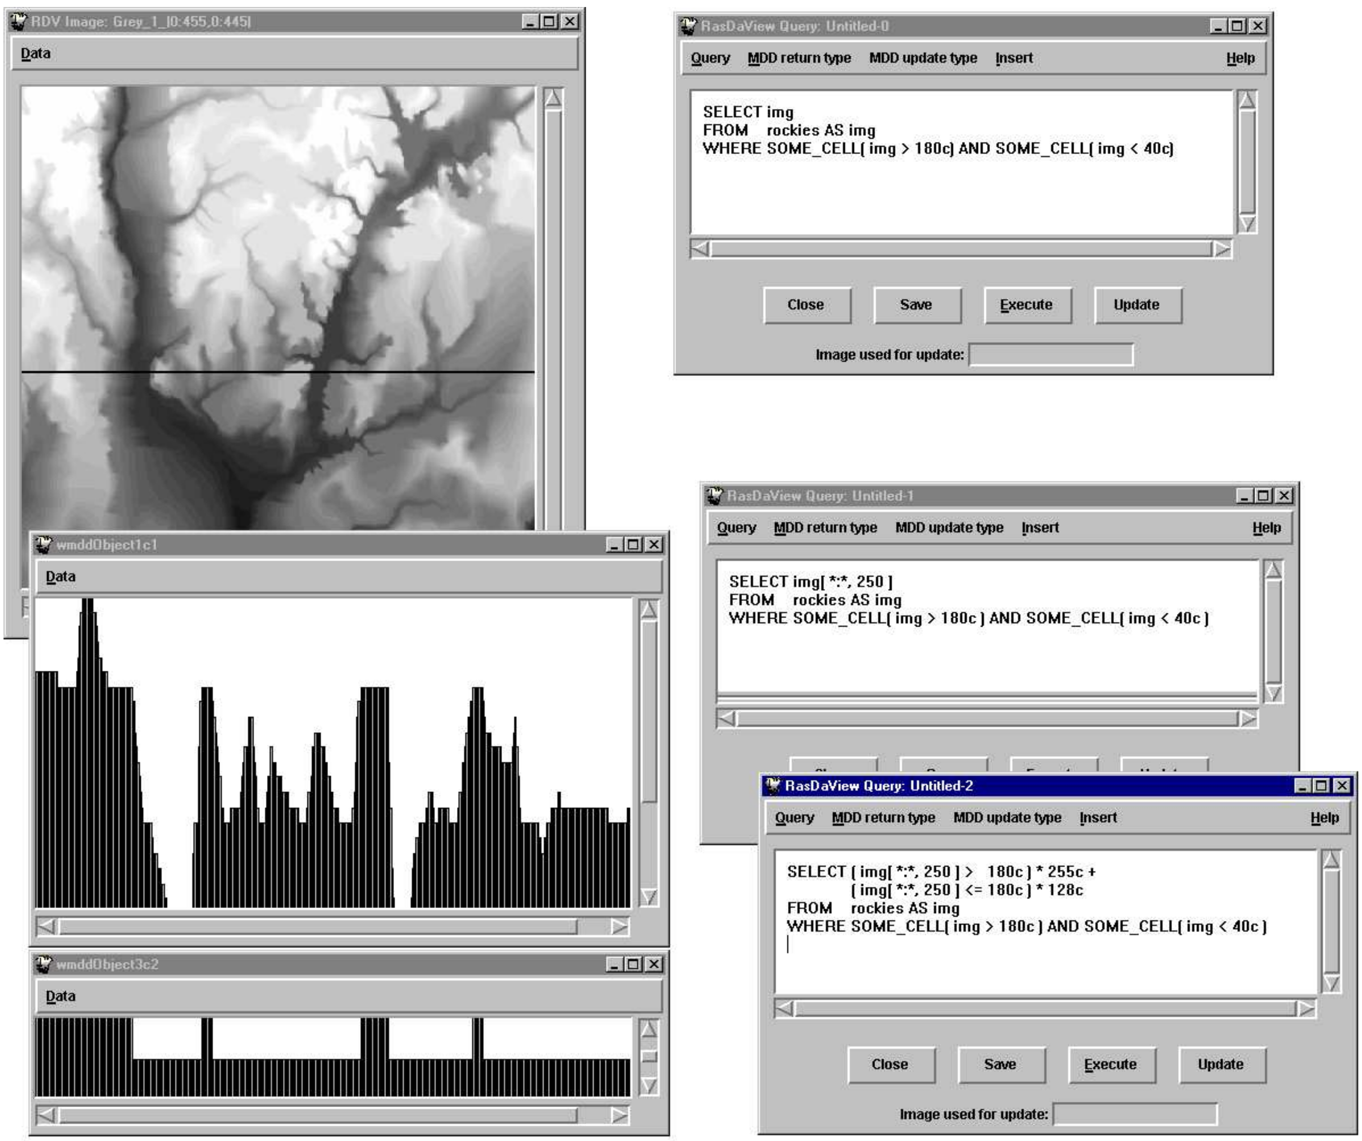

Figure: Sample RasDaMan retrieval using RasDaView for querying and visualization.
From a Digital Elevation Map (DEM) of the Grand Canyon (top left image, corresponding query top right) an East-West section is retrieved and displayed diagrammatically. First, the elevation data themselves are retrieved (middle left, corresponding query middle right), then a threshold-based recoding is done, mapping values less than 180 to 128 and those exceeding 180 to 255 (bottom left, corresponding query bottom right).