

Revolutionary Database Technology for Data Intensive Research

by Martin Kersten and Stefan Manegold

The ability to explore huge digital resources assembled in data warehouses, databases and files, at unprecedented speed, is becoming the driver of progress in science. However, existing database management systems (DBMS) are far from capable of meeting the scientists' requirements. The Database Architectures group at CWI in Amsterdam cooperates with astronomers, seismologists and other domain experts to tackle this challenge by advancing all aspects of database technology. The group's research results are disseminated via its open-source database system, MonetDB.

The heart of a scientific data warehouse is its database system, running on a modern distributed platform, and used for both direct interaction with data gathered from experimental devices and management of the derived knowledge using workflow software. However, most (commercial) DBMS offerings cannot fulfill the demanding needs of scientific data management. They fall short in one or more of the following areas: multi-paradigm data models (including support for arrays), transparent data ingestion from, and seamless integration of, scientific file repositories, complex event processing, and provenance. These topics only scratch the surface of the problem. The state of the art in scientific data exploration can be compared with our daily use of search engines. For a large part, search engines rely on guiding the user from their ill-phrased queries through successive refinement to the information of interest. Limited a priori knowledge is required. The sample answers returned provide guidance to drill down, chasing individual links, or to adjust the query terms.

The situation in scientific databases is more cumbersome than searching for text, because they often contain complex observational data, eg telescope images of the sky, satellite images of the earth, time series or seismograms, and little a priori knowledge exists. The prime challenge is to find models that capture the essence of this data at both a macro- and micro-scale. The answer is in the database, but the 'Nobel-winning query' is still unknown.

Next generation database management engines should provide a much richer repertoire and ease of use experience to cope with the deluge of observational data in a resource-limited setting. Good is good enough as an answer, provided

the journey can be continued as long as the user remains interested.

We envision seven directions of long term research in database technology:

- **Data Vaults.** Scientific data is usually available in self-descriptive file formats as produced by advanced scientific instruments. The need to convert these formats into relational tables and to explicitly load all data into the DBMS forms a major hurdle for database-supported scientific data analysis. Instead, we propose the data vault, a database-attached external file repository. The data vault creates a true symbiosis between a DBMS and existing file-based repositories, and thus provides transparent access to all data kept in the repository through the DBMS's (array-based) query language.
- **Array support.** Scientific data management calls for DBMSs that integrate the genuine scientific data model, multi-dimensional arrays, as first-class citizens next to relational tables, and unified declarative language as symbiosis of relational and linear algebra. Such support needs to be beyond 'alien' extensions that provide user defined functions. Rather, arrays need to become first-class DBMS citizens next to relational tables.
- **One-minute database kernels.** Such a kernel differs from conventional kernels by identifying and avoiding performance degradation by answering queries only partly within strict time bounds. Run the query during a coffee break, look at the result, and continue or abandon the data exploration path.
- **Multi-scale query processing.** Fast exploration of large datasets calls for

partitioning the database based on science interest and resource availability. It extends traditional partitioning schemes by taking into account the areas of users' interest and the statistical stability in samples drawn from the archives.

- **Post-processing result sets.** The often huge results returned should not be thrown at the user directly, but passed through an analytical processing pipeline to condense the information for human consumption. This involves computation intensive data mining techniques and harnessing the power of GPUs in the software stack of a DBMS.
- **Query morphing.** Given the imprecision of the queries, the system should aid in hinting at proximity results using data distributions looked upon during query evaluation. For example, aside from the traditional row set, it may suggest minor changes to the query predicates to obtain non-empty results. The interesting data may be 'just around the corner'.
- **Queries as answers.** Standing on the shoulders of your peers involves keeping track of the queries, their results, and resource requirements. It can be used as advice to modify ill-phrased queries that could run for hours producing meaningless results.

Please contact:

Stefan Manegold

CWI, The Netherlands

E-mail: Stefan.Manegold@cwi.nl