

# O-O, What's Happening to DB2? (Industrial Presentation)

M. Carey, D. Chamberlin,  
S. Narayanan, and B. Vance  
IBM Almaden Research Center  
650 Harry Road, K55/B1  
San Jose, CA 95120-6099

D. Doole, S. Rielau,  
and R. Swagerman  
IBM Toronto Laboratory  
1150 Eglinton Avenue East  
Toronto, Ontario, Canada M3C 1H7

N. Mattos  
IBM Database Technology Institute  
Santa Teresa Laboratory  
555 Bailey Avenue  
San Jose, CA 95141

## ABSTRACT

In this presentation, we will describe a collection of new object-relational features that have been added to IBM's DB2 Universal Database (UDB) system. The features to be described include support for structured types, object references, and hierarchies of typed tables and views. These features will be covered from the perspective of a database designer or end user. In addition to presenting the features presently available in DB2 UDB V5.2, which became available in Fall 1998, we will discuss the expected evolution and impact of this technology over time.

## Keywords

Object-relational database systems, DB2, SQL3, SQL99.

## 1. INTRODUCTION

For over a decade, database researchers have been studying ways to combine the technologies of objects and databases in order to raise the level of abstraction for enterprise information modelers, to reduce the impedance mismatch between the type systems of databases and those of modern programming languages, and to enable databases to support the management of new, possibly complex, user-defined data types. Various approaches have been pursued [1], including the addition of persistence to object-oriented programming languages, the development of a new variety of database systems from the ground up (*i.e.*, object-oriented database systems), the development of object-

oriented middleware to provide an object programming layer on top of relational database systems, and the extension of relational database systems themselves with built-in object capabilities (*i.e.*, object-relational database systems). It is our belief that, for a number of practical reasons, the object-relational approach holds the most promise, at least for general-purpose applications. In particular, the object-relational approach is very attractive because it adds object capabilities to industrially-proven systems that already offer a number of other valuable advanced features (such as highly robust storage managers, mature cost-based query optimizers, server-side support for business logic through constraints and triggers, and scalable parallel query processing engines).

IBM's DB2 Universal Database (UDB) system is following the object-relational approach. DB2 took its first object-relational steps with the 1995 release of DB2 Version 2 for Common Servers, which offered support for user-defined distinct types, user-defined functions, binary and character-based large objects, and triggers (all of which are among the features identified by Stonebraker as elements of object-relational database technology [6]). In 1997, IBM released DB2 Universal Database Version 5.0, which added support for parallel platforms (*e.g.*, the IBM SP2), thereby creating one of the early parallel object-relational database systems [3]. In September 1998, IBM released Version 5.2 of DB2 UDB. Version 5.2 added a collection of significant new object-relational features to DB2; these features emerged from an ongoing object-relational research project at the IBM Almaden Research Center. This presentation will cover those new features of DB2 UDB V5.2 from the perspective of a database designer and/or user of the system.

## 2. OBJECT-RELATIONAL EXTENSIONS

DB2 UDB V5.2 includes new object-relational data definition language (DDL) and data manipulation language (DML) support for the following suite of data modeling features and corresponding query capabilities:

- Hierarchies of user-defined structured types (*i.e.*, named types with attributes) and their subtypes.
- Typed tables, *a.k.a.* "tables of objects", with subtables for holding instances of subtypes.
- Object identifiers and reference types, for directly modeling 1-to-many relationships.
- Typed views, *a.k.a.* "object views", with subviews for dealing with instances of subtypes.
- Semantic extensions to the **SELECT**, **UPDATE**, **INSERT**, and **DELETE** statements of DB2 SQL to support operations on table and view hierarchies.
- An authorization model appropriate for protecting information (including type information) in table and view hierarchies.
- Path expressions for traversing references in SQL queries without having to write joins.
- Additional DML features for inquiring about and/or reacting to a referenced object's runtime type.

More information about these features can be found in a recent IBM technical report [2] that discusses them from both an external user's perspective and a system design and implementation perspective. In addition, the SQL Reference Manual for DB2 UDB V5.2 documents these features in detail and can be obtained online [5].

## 3. INTENDED AUDIENCE

We expect that this presentation will be of interest to SIGMOD attendees for several reasons. First, with the release of UDB V5.2, DB2 has a combination of object-relational features that cannot (at least yet) be otherwise found among products from the major relational database vendors. As of this writing, Oracle does not yet support inheritance or table/view hierarchies, and Informix does not yet support references with path expressions or object views. It should thus be interesting to SIGMOD attendees to see how these object-oriented features look and work together in an object-relational setting. Second, V5.2's

object view support is quite novel, constituting a research contribution in its own right. Object-oriented database researchers are therefore likely to find this aspect of the presentation to be of interest. Finally, IBM has played a major role, working together with several of our object-relational vendor colleagues, in setting the SQL99 standard (which was known until very recently as SQL3 [4]) in the area of object extensions for the SQL language. Learning about the new object-relational features of V5.2 is thus a way for SIGMOD attendees to get a concrete sense of how SQL99 has shaped up in this area.

The plan for this presentation is to combine viewgraphs together with a real-time demonstration of the system. The presentation will be technical, covering the new object-relational features in DB2 UDB V5.2 from a SQL DDL and DML language perspective. Viewgraphs will be used to introduce each of the new features; each feature will also be demonstrated in real-time (using DB2 UDB on NT running on an IBM ThinkPad) on a sample database during the talk. The talk will close with a recap of what UDB V5.2 offers today and a general discussion of future directions for IBM's object-relational DB2 work in the V6/V7 timeframe and beyond.

## 4. REFERENCES

- [1] Carey, M., and DeWitt, D., "Of Objects and Databases: A Decade of Turmoil", *Proc. of the 1996 International Conference on Very Large Data Bases*, Mumbai (Bombay), India, September 1996.
- [2] Carey, M., Chamberlin, D., Narayanan, S., Vance, B., Doole, D., Rielau, S., Swagerman, R., and Mattos, N., *O-O, What Have They Done to DB2?*, Research Report No. RJ 10132, IBM Almaden Research Center, October 1998.
- [3] Chamberlin, D., *A Complete Guide to DB2 Universal Database*, Morgan Kaufmann Publishers, San Francisco, CA, 1998.
- [4] Eisenberg, A., and Melton, J., "SQL:1999, formerly known as SQL3", *ACM SIGMOD Record* 28(1), March 1999.
- [5] IBM DB2 Product and Service Technical Library, <http://www.software.ibm.com/data/db2/library/>.
- [6] Stonebraker, M. *Object-Relational Database Systems: The Next Great Wave*, Morgan Kaufmann Publishers, San Francisco, CA, 1996.