CALL FOR PAPERS
Information Systems: Special Issue on Data Warehousing

Editors

Patrick O'Neil
University of Massachusetts/Boston
poneil@cs.umb.edu

Kenneth A. Ross
Columbia University
kar@cs.columbia.edu

Introduction. A "Data Warehouse" is a database system that provides query capability to analysts engaged in extracting strategic information about a (usually commercial) enterprise. Data Warehouses can use a number of different data models and user interfaces, including SQL queries on Star Schemas, OLAP (ROLAP and MOLAP), and Data Mining, where detail queries are performed automatically for high-level goals set by the analyst. In most cases, the data about the enterprise held in the Warehouse has been extracted from the enterprise's operational (OLTP, Point-Of-Sale) system, and processed in various ways (brought to a new data model, cleansed, indexed, preaggregated) to make queries more efficient and responsive to analyst's needs. Data will often be brought in from other sources as well.

A Data Warehousing system at an enterprise can have multiple goals. One goal can be to answer the rather fundamental concerns of senior management: "We sell various products in various markets and measure our progress over time; how are we doing?" More complex questions arise from this that can provide strategic advantage: "To what can we attribute variations in sales by product, market, and time period (as well as by placement on shelves, packaging, offered promotions, etc.)? How can we build on this understanding to increase our total sales?" Other goals include providing suppliers with good information for reordering, detecting bottlenecks in complex supply and demand chains, and understanding customer requirements.

Details of the Call. Because of the widespread research activity in this topic, Information Systems is planning a special issue on Data Warehousing Systems for Winter 2000/1. Our goal is to collect papers with important new insights and experiences, and put together a strong issue. Theoretical papers should include a solid motivation for why the stated results are applicable in real applications. Applied papers that give insightful descriptions of working systems in actual use are particularly solicited. Topics of interest include, but are not limited to, the following.

- Novel Indexing for Performance
- Novel query processing algorithms
- Materialized views
- MOLAP, ROLAP, other models
- Data Mining
- Data Cleansing
- Data Loading
- Updating (Especially with following two topics)
- 24X7 Availability
- Stability (Repeatability of Queries)
- Schema Design

Papers should be at most 30 pages long, double-spaced, in font size 10 or larger with one-inch margins on all sides. Five copies of each paper should be sent to: Patrick O'Neil; Dept., of Math. and C.S.; UMass/Boston; Boston, MA 02125-3393; USA. Alternatively, a paper in PDF format can be sent to poneil@cs.umb.edu by email. Note the following important dates:

Date	Milestone
04/01/2000	Deadline for submission of papers
07/01/2000	Notification of 1st-round referee results
09/01/2000	Deadline for submission of revised papers
10/15/2000	Notification of final acceptance/rejection
11/15/2000	Deadline for having final papers at the publisher