

Go Faster!

The TransRelational™ Approach to DBMS Implementation

by

C. J. Date

ABOUT THIS SEMINAR

In the field of scientific endeavor, an idea emerges from time to time that is so startlingly novel, and so dramatically better than anything that went before, that it can truly be described as a breakthrough. The relational model provides the obvious example in the database world; almost everything done in that world since the relational model came along stands as testament to the radical nature and impact of that one brilliant idea. And now we are witnessing the birth of what looks set to be another major breakthrough: **The TransRelational™ Model**. In this speaker's opinion, the TransRelational model is likely to prove the most significant advance in this field since Codd gave us the relational model, nearly 35 years ago.

So what is the TransRelational model? In essence, it is an implementation technology; it represents among other things a radically new and elegant approach to the question of implementing the relational model. When the relational model first appeared, skeptics accused it of being impossible to implement efficiently. Over the next few decades, therefore, vast sums of money were spent on attempts to prove the skeptics wrong. And those efforts did not prove in vain; today's products do perform fairly well, if not always spectacularly. But at what cost? The products have become quite notoriously unwieldy and difficult to manage, with their huge array of storage structures, access methods, optimization techniques, tuning knobs, analysis utilities, performance hints, installation options, and so on, and the job of the database administrator has become virtually impossible.

The TransRelational model has the potential to change all that. It provides a totally new approach to implementation, one that is dramatically different from those that have been tried in the past and found wanting (including *all* of the approaches encountered in today's SQL products). Such an implementation would be orders of magnitude faster than, and would deliver a far greater degree of data independence than, today's SQL products. It would also greatly simplify the job of the DBA.

Here are a few more specific features of the TransRelational approach:

- The data is effectively stored in many different sort orders at the same time.
- The data occupies a fraction of the space required for a conventional database today.
- Indexes and other conventional access paths are completely unnecessary.
- Optimization is much simpler than it is with conventional systems; often, there is just one obviously best way to implement any given relational operation. In particular, the need for cost-based optimizing is eliminated.
- There is no need to compile database requests ahead of time for performance.
- Join performance is linear!—meaning, in effect, that the time it takes to join twenty relations is only twice the time it takes to join ten (loosely speaking). It also means that joining twenty relations, if necessary, is feasible in the first place; in other words, the system is *scalable*.
- Logical design can be done properly (in particular, there is never any need to "denormalize for performance").
- Physical database design can be completely automated.
- Database reorganization as conventionally understood is completely unnecessary.
- In general, far fewer human decisions are needed.

In a nutshell, the TransRelational model allows us to build DBMSs that—at last!—truly deliver on the full promise of the relational model.

TOPIC OUTLINE

Part 1:

- Preliminaries
- The historical context
- Three levels of abstraction
- Core concepts
- Core concepts (cont.)

Part 2:

- Major-to-minor orderings
- Condensed columns
- Merged columns
- Relational operators
- Review and analysis

DURATION

Part 1: 2½-3 hours (approx.); Part 2: 2½-3 hours (approx.).

WHO SHOULD ATTEND

Anyone interested in the field of data management who wants to learn about a truly dramatic new development in that field should benefit from this seminar. The only prerequisites are an elementary appreciation of data management concepts, techniques, and issues, including some knowledge of the relational model in particular.

OBJECTIVES

On completion of this seminar, attendees will:

- Understand what's wrong with conventional DBMS implementation technology
- Understand how the TransRelational Model is dramatically different—what makes it "blindingly fast," how it greatly simplifies the job of the DBA, and why it holds hope for the future
- Be in a good position to make a transition to that technology when it arrives

DOCUMENTATION

Attendees will receive a workbook containing copies of the speaker's slides and background reading material.

SPEAKER: Chris Date

C. J. Date is an independent author, lecturer, researcher, and consultant, specializing in relational database technology. He is best known for his book *An Introduction to Database Systems* (eighth edition, Addison-Wesley, 2004), which has sold some 725,000 copies and is used by several hundred colleges and universities worldwide. He is also the author of many other books on database management, including most recently:

- From Morgan Kaufmann: *Temporal Data and the Relational Model* (coauthored with Hugh Darwen and Nikos A. Lorentzos, 2003)
- From O'Reilly: *Database in Depth: Relational Theory for Practitioners* (2005)
- From Addison-Wesley: *Databases, Types, and the Relational Model: The Third Manifesto* (coauthored with Hugh Darwen, to appear 2006)

Another book, *Go Faster! The TransRelationaltm Approach to DBMS Implementation*, is also due for publication in the near future.

Mr. Date enjoys a reputation that is second to none for his ability to communicate complex technical subjects in a clear and understandable fashion.

