

Index Research: Forest or Trees?

Joseph M. Hellerstein
University of California, Berkeley
Berkeley, CA 94720-1776
jmh@cs.berkeley.edu

Indexes and access methods have been a staple of database research – and indeed of computer science in general – for decades. A glance at the contents of this year’s SIGMOD and PODS proceedings shows another bumper crop of indexing papers.

Given the hundreds of indexing papers published in the database literature, a pause for reflection seems in order. From a scientific perspective, it is natural to ask why definitive indexing solutions have eluded us for so many years. What is the grand challenge in indexing? What basic complexities or intricacies underlie this large body of work? What would constitute a successful completion of this research agenda, and what steps will best move us in that direction? Or is it the case that the problem space branches in so many ways that we should expect to continuously need to solve variants of the indexing problem?

From the practitioner’s perspective, the proliferation of indexing solutions in the literature may be more confusing than helpful. Comprehensively evaluating the research to date is a near-impossible task. An evaluation has to include both functionality (applicability to the practitioner’s problem, integration with other data management services like buffer management, query processing and transactions) as well as performance for the practitioner’s workloads. Unfortunately, there are

no standard benchmarks for advanced indexing problems, and there has been relatively little work on methodologies for index experimentation and customization. How should the research community promote technology transfer in this area? Are the new extensibility interfaces in object-relational DBMSs conducive to this effort?

Panelists

The members of this panel provide a broad mix of research perspectives and real-world experience:

- Hans-Peter Kriegel, University of Munich
- Christos Papadimitriou, UC Berkeley
- David Lomet, Microsoft Research
- Christos Faloutsos, CMU
- Raghu Ramakrishnan, University of Wisconsin
- Paul Brown, Informix Corporation

Permission to make digital or hard copies of part or all of this work or personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

MOD 2000, Dallas, TX USA

© ACM 2000 1-58113-218-2/00/05 . . . \$5.00