Reminiscences on Influential Papers

Kenneth A. Ross, editor

I continue to invite unsolicited contributions. See `http://www.acm.org/sigmod/record/author.html` for submission guidelines.

Amr El Abbadi, University of California, Santa Barbara, `amr@cs.ucsb.edu`.

[G. Wuu and A.J. Bernstein. Efficient Solutions to the Replicated Log and Dictionary Problems. PODC 1984: 233-242.]

This is a beautifully simple paper that I feel encompasses many ideas that keep reappearing in different guises every decade or so! The paper proposes the replication of a dictionary (basically a set of key and value pairs) to all relevant sites in a distributed system. Updates and deletes are propagated in a lazy manner through the system as sites communicate with each other using a simple notion of a log. Queries are answered based on the local copy. The notion of correctness is based on the causal dependency between operations. A simple data structure keeps track of "who knows what", and is used for reducing the propagated information as well as garbage collection.

The framework is elegant and the solution is simple. I always enjoy teaching it in class because it is quite easy to explain and yet can be used as a platform to address many foundational issues. It has been the foundation for much work since the mid-eighties; of particular note is the Bayou system at Xerox PARC. I also view it as being significant in providing an alternative framework to the standard serializability notion of correctness. More recently, it can also be regarded as the foundation of a peer-to-peer approach for maintaining distributed information. Also, it represents the original closeness of the distributed systems and the database communities, which, unfortunately is happening less and less these days. The paper appeared in the Principles of Distributed Computing conference in 1984. The original formulation of the dictionary problem was presented by Fischer and Michael in our own Principles of Database Systems conference in 1982. At the time there was much cross fertilization between the 2 communities. This is essentially a data management problem, which was solved in a peer-to-peer setting. In the early eighties many influential results were being developed by both communities, e.g., a little known fact is that one of the most important results in distributed computing, namely the impossibility of consensus in a distributed system with one fault, by Fischer, Lynch and Paterson also appeared in PODS in 1983. The two communities share many similar problems, and I would like to take this opportunity to advocate more such cross fertilization, which will be beneficial to all of us.

Theodore Johnson, AT&T Labs Research, `johnsont@research.att.com`.

[Gennady Antoshenkov. Byte-aligned Bitmap Compression. Technical Report, Oracle, 1994 / U.S. Patent 5,363,098.]

In the late '90s, I was considering the problem of indexing very large data warehouses (AT&T Labs-Research has a lot of them). Bitmaps seemed to be ideally suited for many applications, but their scalability was a weakness (one bitmap is required for each unique value of the indexed attribute). The bitmaps can be compressed to achieve competitive storage requirements, but most bitmap compression/decompression algorithms are slow.

Pat O'Neil gave me a copy of Gennady's techical report. Using fundamentally simple (and therefore fast) mechanisms, his bitmap compression scheme automatically shifts between an uncompressed and a compressed form as appropriate. With this simple adaptivity, his bitmap encoding scheme achieves both high compression and fast Boolean operations, making bitmap indices a practical choice in a wide variety of applications. Many

commercial implementations of compressed bitmaps use these ideas, and my work on bitmap indexing was motivated by Antoshenkov's work. This paper is interesting for two additional reasons. First, it is a simple and effective solution because it approached the problem in a new way. Second, it is published as a United States patent, and I suspect that a great deal of industrial database research is also.

Richard T. Snodgrass, University of Arizona, `rts@cs.arizona.edu`.

[James Clifford, David Warren. Formal Semantics for Time in Databases. ACM TODS 8(2), June 1983, pages 214–254. (Extended abstract in XP1 Workshop on Database Theory, 1981.)]

Way back in 1983 I was a wet-behind-the-ears assistant professor with a freshly completed dissertation on distributed operating systems when I came across Jim Clifford's dissertation work in TODS coauthored with his advisor. This beautiful paper impressed upon me the possibility that temporal databases (termed *historical databases* in that paper; the scope was later broadened with the more accurate term to include future data) was a relevant subject of study, with a pleasing subtlety and interesting fundamental questions and that time was an *orthogonal* aspect of data, rather than just another column in a table. I was extremely fortunate to have read this paper at a highly impressionable age, for it literally changed my life and sent me on a delightful journey of two decades that continues into temporal databases.

This paper opened up many research questions, such as whether to timestamp attributes or tuples, how to specify temporal data dependencies and temporal constraints, how to state temporal queries in a simple, convenient, yet adequately expressive manner, and how to efficiently store time-varying data, especially when the data model admitted much redundancy. In raising these issues, the paper consistently argued for providing a formal model-theoretic semantics. We now know the answers to these questions, in part because Jim raised them in the first place.

The paper was also influential to me in another way. It indirectly, by example, encouraged adopting insights from computational linguistics. I now feel strongly that linguistics offers a useful view into the mind of a database designer or application programmer. The way people think is often mirrored in the way that they express things in language, and so linguistic structures provide valuable clues on how humans structure their understanding of the world. By striving to make our formalisms consistent with language, we can make these formalisms consistent with our internal thought processes.

Jim passed away unexpectedly in 1995; the SIGMOD conference proceedings that year was dedicated to him and to Paris Kanellakis. The tragedy was not that someone who had produced so many ideas and great papers had died, but rather that all of us were thus deprived of the insights that Jim could have provided in the years that followed, and also that younger researchers were deprived of the joy knowing this wonderful person.
