

In the past five years, columnar storage technologies have gained momentum in the analytical marketplace. This resulted from (i) a number of new successful entrants in the database market with products based on columnar technology, as well as (ii) established vendors introducing columnar products or even integrating columnar technology deep inside their existing products.

In this issue, you find a number of papers that describe systems that fit one or even both of these categories: Infobright and MonetDB as pure column stores; and Virtuoso, Vectorwise, IBM's Blink based products, Microsoft SQLserver, and SAP HANA (through its P*TIME subsystem) are products that combine row- and column-technology; this also holds for the research systems HYRISE and Hyper. Given that column store systems have been on the market now for a few years, the commercial system papers also describe customer reactions, so we can assess how column stores are holding up in practice. These papers contain concise system descriptions that hopefully will inspire the reader, and also form an entry point for further reading.

Previous to these five years of major commercial adoption, there has been a long research track in column stores, most visibly in the MonetDB system, which is completing its second decade of research history. MonetDB has proven to be a powerful platform for new interesting systems research such as on-the-fly indexing ("cracking") and adaptive result caching ("recycling"). But not only at CWI does columnar technology inspire new academic research in new topics, evidenced by descriptions of HYRISE and Hyper. Both latter papers address, in different ways, the issue of combining row and columnar formats, among other topics.

While the success of specialized columnar systems seemed to underline the end of the "one system fits all" paradigm as proclaimed by Michael Stonebraker, this issue clearly shows that this is still a debatable proposition. Both the Microsoft SQLserver as well as the Openlink Virtuoso systems show that tight integration of columnar technology in row-based systems is both possible and desirable. Both systems are deeply integrated, as they do not stop at only superficially adopting columnar storage, but also vectorized large parts of their execution systems to reap its query processing benefits. Though Virtuoso probably is the lesser well-known system (database practitioners working with RDF will surely know it), it is an especially interesting case, as the upcoming version 7 described and microbenchmarked here integrates columnar and vectorized technology fully throughout data storage, query execution and event MPP cluster infrastructure. In this system, rows and columns fully co-exist, which enables interesting apples-to-apples comparisons.

The idea that one would have to choose between row- or columnar-systems is also contradicted by the work in HYRISE and Vectorwise, where even during execution tuples are represented partly columnar and partly row-wise – this also true in Blink due to packing of columns in machine words. The Hyper paper confronts "one system fits all" head on, arguing for the contrary. It shows in proof-of-concept that by machine-supported forms of transaction isolation, and efficient query compilation techniques, one system can both compete or exceed the best specialized alternatives in OLTP and OLAP workloads.

In all, columnar systems research leads to interesting questions and will continue to influence future developments in database architecture.

Let me hereby thank all authors for their efforts, and express my hope that you will enjoy reading the resulting material.