DOI:10.1145/2380656.2380659

Why Open Access?

[redacted]

be willing to pay more simply because [redacted]

Not the Database World We Know

Communications readers have a right to expect accuracy. Sadly, accuracy is not always what they get. The article "All Your Database Are Belong to Us" by Erik Meijer (Sept. 2012) contains so many inaccuracies, confusions, and errors regarding "the database world" it is difficult to read coherently. The first paragraphs alone contain more egregious misstatements than most entire articles or papers. For the record: "The raw physical data model" is categorically not "at the center of the [relational database] universe." Queries do not "assume intimate details of the data representation (indexes, statistics, metadata)." While database technology relies on "The Closed World Assumption," this assumption has nothing to do with what the author apparently meant. Every phrase in "Exposing naked data and relying on declarative magic becomes a liability" relies on at least one counterfactual. "Objects should hide their private data representation, exposing it only via well-defined behavioral interfaces." But this is exactly what the relational model does except (unlike OO) it adopts an interface discipline that makes ad hoc query and the like possible. "In the realm of [data] modelers, there is no notion of data abstraction." Astoundingly wrong. "[Database technology necessarily involves] a computational model with a limited set of operations." False. Although the (very powerful, well-defined, provably correct) required set of relational operations is small, the sky's the limit on derived relational operations or operations that define abstract data type/domain behavior.

The author's unfounded antipathy toward relational databases dominates even his application of CAP: "The problem with SQL databases... is the assumption that the data... meets a bunch of consistency constraints that is difficult to maintain in an open ['anything goes'?] distributed world." CAP does not eliminate this requirement; "...the hidden cost of forfeiting [system-enforced] consistency... is the need [for the programmer] to know the system's invariants."[1] Nor can programmers "...design their systems to be robust...to inconsistency." Once data inconsistency invades a computationally complete system, it is not even, in general, detectable, and all bets are off. Consistency must be enforced, hence constraints. The author seemed to equate detecting abnormal execution with enforcing logical data consistency. No wonder confusion abounds; CAP consistency is single-copy consistency, a subset of what ACID databases provide, yet the Gilbert/Lynch CAP proof relies on linearizability, a more stringent requirement than the serializability ACID databases need or use.

And so on... Deconstructing the entire article properly would take more time than we care to devote, but the foregoing should suffice to demonstrate its fallaciousness. We hope the author is not teaching these confusions, errors, logical inconsistencies, and fallacies.

It is difficult even to believe the article was peer reviewed. Indeed, it is truly distressing it did not demonstrate even minimal understanding of one of the most important contributions to computing: the relational model. We can only deplore *Communications*'s role in promulgating such a lack of understanding.

C.J. Date, Healdsburg, CA,
and **D. McGoveran**, Boulder Creek, CA

Reference

1. Brewer, E. CAP twelve years later: How the 'rules' have changed. *IEEE Computer 45*, 2 (Feb. 2012), 23–29.

Author's Response:

The purpose of the article was not to criticize the relational model but to point out how building industrial-strength systems using today's relational database systems requires leaving the ivory tower and dealing with a morass of ad hoc extensions to the clean mathematical basis of first-order predicate logic. Rather than depend on pure sets and relations, developers need to think in terms of (un)ordered multisets. For the sake of efficiency and lock-contention avoidance, transactions allow for various isolation levels that clearly violate the ACID guarantees of Platonic transactions. The article also considered whether in the new world of the Cloud we should view as complementary computational models that fundamentally address loosely coupled distributed systems, like Carl Hewitt's Actors.

Erik Meijer, Delft, The Netherlands

Communications welcomes your opinion. To submit a Letter to the Editor, please limit yourself to 500 words or less, and send to letters@cacm.acm.org.