

Before anybody get excited about map/reduce based database systems, it's important to understand their limitations.

For those unfamiliar with the technology, map/reduce is a technique pioneered by Google to exploit parallelism across a large number of potentially unreliable servers. A problem is first broken down in a finite number of independent tasks (the map phase), which are parceled out to different servers for execution. The individual results are then integrated (the reduce phase) and a single answer returned. The reduce phase also restarts any failed tasks.

Both Google (BigTable) and Amazon (SimpleDB) have built what they call databases on map/reduce.

The essential limitation of the technology that it is difficult, probably impossible, to support transactions since, after all, the various servers are independent. To compensate, both BigTable and SimpleDB use complex rows that are essentially self defining and support things like versions and repeating groups. Updates in both systems are atomic but operate on single rows. This is a data model that works very well for shopping carts and very poorly for almost everything else.

Map/reduce is very good technology for a large number of computationally large problems. But I doubt that it has much, if anything, to offer to database systems.

--

Jim Starkey  
Founder, NimbusDB, Inc.  
978 526-1376

Quoted from: Firebird-Architect@yahoogroups.com