

1077-05-1447

Jeremy Kepner*, (kepner@ll.mit.edu), MIT Lincoln Laboratory, Lexington, MA. *Spreadsheets, Big Tables, and the Algebra of Associative Arrays.*

Spreadsheets are used by nearly 100M people every day and may be the most commonly used analytical structure on Earth. Likewise triple stores (e.g., Big Table, Dynamo, and HBase) store a large fraction of the analyzed data in the world and are the backbone of modern web companies such as Google, Amazon, and Yahoo. Both spreadsheets and big tables can hold diverse data (e.g., strings, dates, integers, and reals) and lend themselves to diverse representations (e.g., matrices, functions, hash tables, and databases). Triple stores are highly scalable and run on commodity clusters, but lack interfaces to support efficient development of mathematical analytics. D4M (Dynamic Distributed Dimensional Data Model) has been developed to provide a mathematically rich interface to triple stores (and structured query language “SQL” databases). D4M allows linear algebra to be applied to databases. Using D4M, it is possible to create highly complex composable analytics with significantly less effort than using traditional approaches. The central mathematical concept of D4M is the “associative array” which combines spreadsheets, triple stores and sparse linear algebra. This talk describes the D4M technology and the group theory and fuzzy algebra foundations of associative arrays. (Received September 19, 2011)