



Contribution ID: 225

Type: **Oral presentation to parallel session**

The Drillbit column store

Thursday, 17 October 2013 13:53 (20 minutes)

In comparison to storing data packed by event, column data stores store event variables or sets of event variables in individual data packs. One well-known example is the CERN ROOT library's TTree, which has a mode where it behaves like a column store. Columnar data stores can offer fast processing of a subset of the event structure or individual variables.

In the experimental Drillbit column store we explore the encoding of Google protocol buffer data structures into columns, using a method used in the internal Google Dremel architecture. In addition, Drillbit aims to provide a robust mechanism to synchronize event variables stored in different files, providing a guarantee to the analyst that the event or partial event has been reassembled correctly. By using blockwise unique identifiers and enforcing event ordering in blocks of events, the performance problems usually associated with database joins are avoided. For reduced analysis datasets, the Drillbit data structure allows efficient removal of events, object variables or subsets of objects, even while keeping the full alignment and compatibility with non-reduced datasets at all levels.

Preliminary studies on real-life ROOT analysis datasets have yielded exciting results, indicating a possible gain of about a quarter in storage space while using the same compression algorithm and settings. In addition, an experimental analysis library which is compatible with a subset of the TTree API showed performance on par with or exceeding the ROOT TTree.

Finally, in connection with an in-development dynamic event model, Drillbit could make it practical to do more cache-efficient computations on small numbers of variables, as well as providing several opportunities to use multiple cores. For analysts, Drillbit could allow fast and reliable retrieval of only the relevant analysis variables, and a simple way to share new data corrections and analysis objects.

Primary authors: Dr EBKE, Johannes (TNG Technology Consulting); Mr WALLER, Peter (University of Liverpool (GB))

Presenter: Dr EBKE, Johannes (TNG Technology Consulting)

Session Classification: Data Stores, Data Bases, and Storage Systems

Track Classification: Data Stores, Data Bases, and Storage Systems