

Lesson 4

Column family, RC, ORC, Parquet and Tabular Data Stores

Columnar Data Store

- A way to implement a schema is the divisions into columns
- Storage of each column, successive values is at the successive memory addresses

Columnar Data Store

- Analytics processing (AP) In-memory uses columnar storage in memory
- A pair of row-head and column-head is a key-pair. The pair accesses a field in the table.

Column Data Store

- All values in successive fields in a column consisting of multiple rows save at consecutive memory addresses. This enables fast accesses during in-memory analytics, which includes CPU accesses and analyses

Column family Big-Data

- A storage in logical groups of column families
- May be similar to columns of sparse matrix
- Use a pair of row and column keys to access fields in the columns

Examples of Column Family Data Store

- HBase, BigTable, HyperTable and Cassandra
- The following example explains a column-family data store and why On-line Analytics Processing (OLAP) is fast in-memory column data store in memory:

Figure 3.6 Fields in columnar storage and addresses in memory

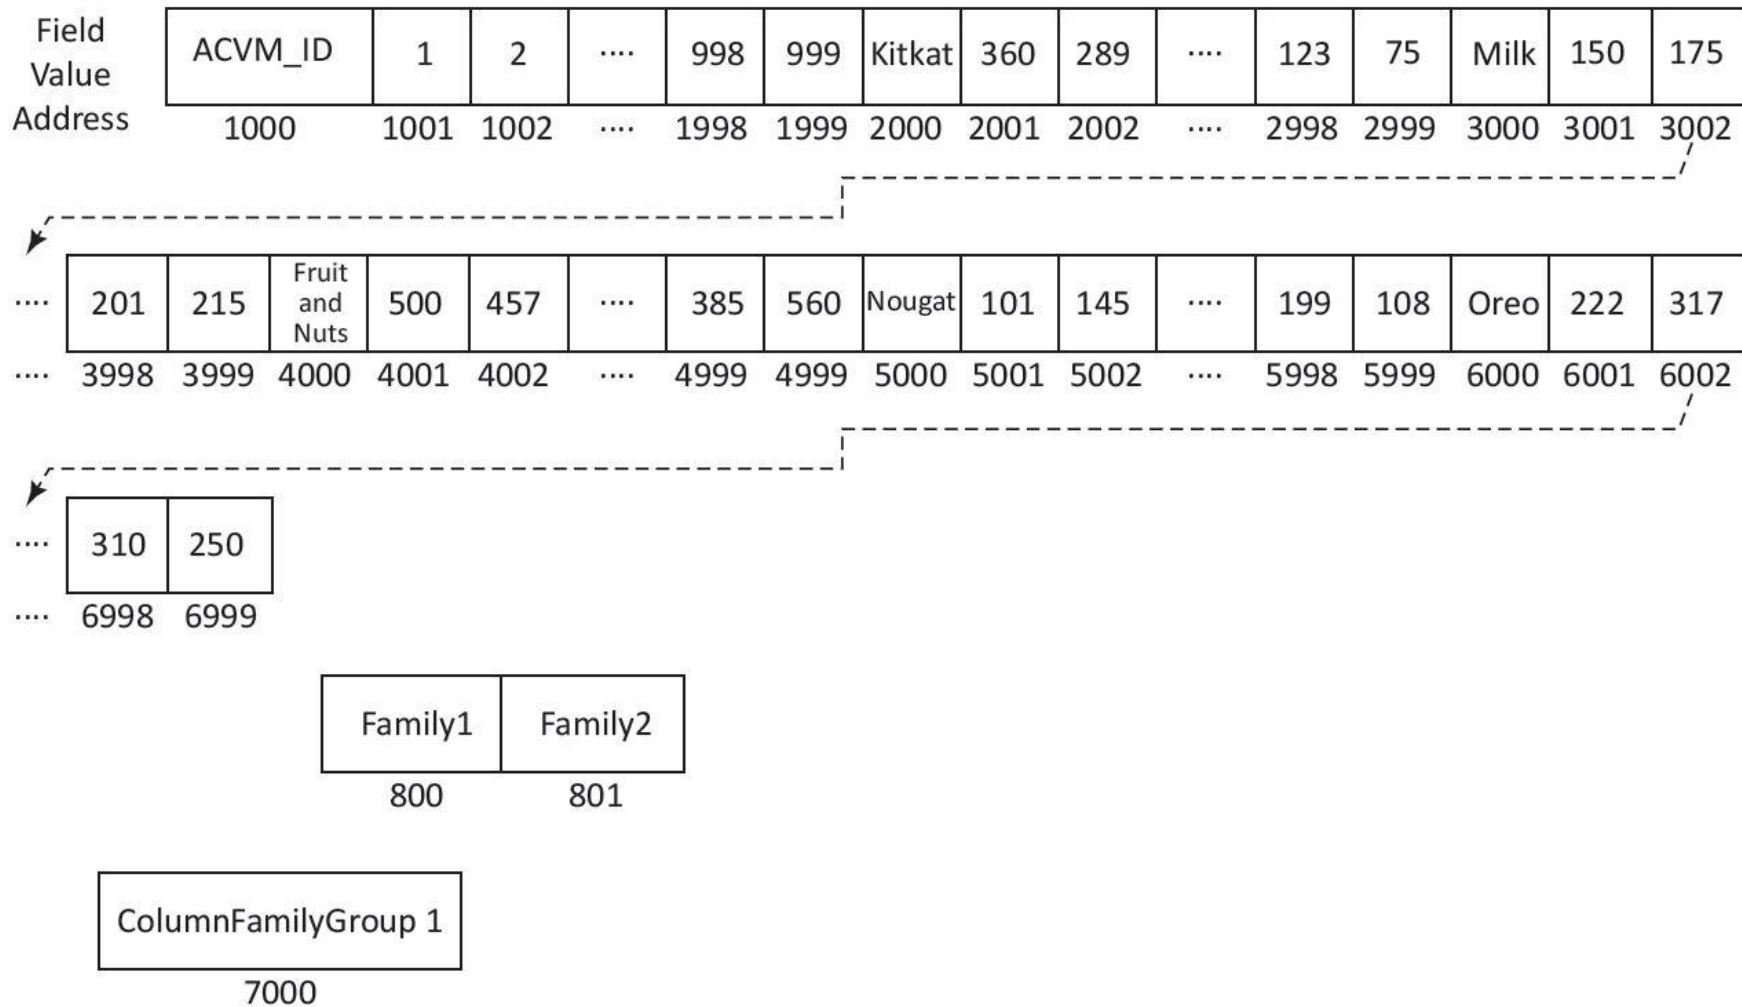


Table 3.3 Each day's sales of chocolates on 999 ACVMs

- Columns in Table 3.3 for KitKat and Milk form a group as one family. Columns for Fruit and Nuts, Nougat, and Oreo form a group as second family

Table 3.3 Each day's sales of chocolates on 999 ACVMs

- The key for one family is 'Popular Flavours Family' and second family is 'Costly Flavours Family'. The keys of column families can save at the addresses 800, 801,

Sparse Column Fields

- A row may associate a large number of columns but contains values in few column fields
- Similarly, many column fields may not have data. Columns are logically grouped into column families. Column-family data stores are then similar to sparse matrix data

Characteristics of Columnar Family Data Store

- Very high performance and scalability,
- Moderate level of flexibility and
- Lower complexity when compared to the object and graph databases
- No optimization for Join

Advantages of column stores are:

- Scalability:
- Partitionability:
- Availability
- Tree-like columnar structure
- Adding new data at ease
- Querying all the field values in a column in a family

BigTable Data Store

1. Massively scalable NoSQL BigTable scales up to 100s of petabytes.
2. Integrates easily with Hadoop and Hadoop compatible systems.

BigTable Data Store

- Example 3.7:
- Column fields which have keys to access a field not only by row ID and Column ID but also include the timestamp and attributes in a row
- Example shows the column-keys for accessing column fields of a column .

Hive Columnar (RC) file-format records

- Used for querying
- RC is the best choice for intermediate tables for fast column-family store in HDFS with Hive
- Serializability of RC table column data is the advantage. RC file is DeSerializable into column data.

Example 3.8

- The values in each column are the records in file for each row group. Each row-group data is like a column of records which stores in the RC file.

ORC File Formats

- An ORC (Optimized Row Columnar) file consists of row-group data called stripes
- ORC enables concurrent reads of the same file using separate RecordReaders.
- Metadata store uses Protocol Buffers for addition and removal of fields.¹

ORC File Formats

- ORC is an intelligent Big Data file format for HDFS and Hive.
- An ORC file stores a collection of rows as a row-group. Each row-group data store in columnar format. This enables parallel processing of multiple row-groups in an HDFS cluster.

ORC File Formats

- ORC is an intelligent Big Data file format for HDFS and Hive.
- An ORC file stores a collections of rows as a row-group.
- Each row-group data store in columnar format. This enables parallel processing of multiple row-groups in an HDFS cluster.

Parquet nested hierarchical columnar-storage

- Nesting sequence is the table, row group, column chunk and chunk page. Apache Parquet file is columnar-family store file
- .

Parquet nested hierarchical columnar-storage

- Apache Spark SQL executes user defined functions (UDFs) which query the Parquet file columns (Section 5.2.1.3). A programmer writes the codes for an UDF and creates the processing function for big long queries.

Parquet nested hierarchical columnar-storage

- A programmer writes the codes for an UDF and creates the processing function for big long queries.

Summary

We learnt :

- Features of Columnar Data Store
- BigTable
- RC
- ORC and
- Parquet data stores

End of Lesson 4 on
**Column family, RC, ORC,
Parquet and Tabular Data Stores**