

Data Model Exercise

You should read the BigTable paper (possibly watch the presentation) - URLs are on the web site. You can also look up the H-Base documentation for more information about the client API: http://hadoop.apache.org/hbase/docs/current/api/org/apache/hadoop/hbase/client/package-summary.html#package_description. H-Base is the Apache implementation of BigTable: <http://hadoop.apache.org/hbase/>.

Consider a 2-dimensional array d2 of 4x4 integers, and a 3-dimensional array d3 of 4x4x4 integers:
In pseudo-java: `int [4] [4] d2; int [4] [4] [4] d3;`

Question 1: How would you represent d2 and d3 in SQL? How would you represent a query to access the current version of d2[2][3]? What happens if d3[2][1][0] is deleted? Updated? Can you retrieve it?

Question 2: How would you represent d2 and d3 in BigTable (H-Base)? How would you represent a query to access the current version of d2[2][3]? What happens if d3[2][1][0] is deleted? Updated? Can you retrieve it?