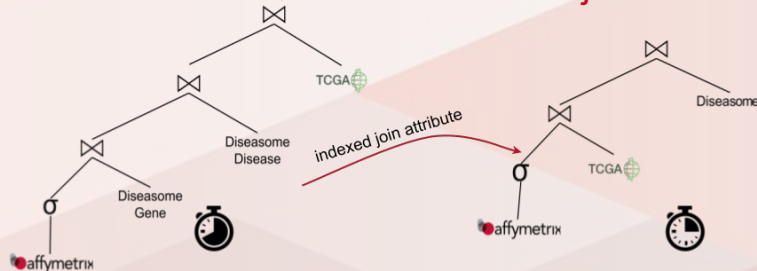


Optimizing Federated Queries Based on the Physical Design of a Data Lake

by Philipp D. Rohde and Maria-Esther Vidal


- find efficient and effective query execution plans
- consider data and infrastructure
- maximize answer completeness
- minimize query execution time

Heuristic 1: Push down joins



Push down joins between two sub-queries if both sub-queries can be evaluated in the same data source and the join attribute is indexed.

Heuristic 2: Push up instantiations



Perform filters at query engine level unless the filtered attribute is indexed and the network speed is slow.