



Optimizability of OGC Standards Implementations - a Case Study

D. Misev and P. Baumann

Jacobs University Bremen gGmbH, EECS, Bremen, Germany (p.baumann@jacobs-university.de)

Why do we shop at Amazon? Because they have a unique offering that is nowhere else available? Certainly not. Rather, Amazon offers (i) simple, yet effective search; (ii) very simple payment; (iii) extremely rapid delivery. This is how scientific services will be distinguished in future: not for their data holding (there will be manifold choice), but for their service quality. We are facing the transition from data stewardship to service stewardship.

One of the OGC standards which particularly enables flexible retrieval is the Web Coverage Processing Service (WCPS). It defines a high-level query language on large, multi-dimensional raster data, such as 1D timeseries, 2D EO imagery, 3D x/y/t image time series and x/y/z geophysical data, 4D x/y/z/t climate and ocean data. We have implemented WCPS based on an Array Database Management System, rasdaman, which is available in open source.

In this demonstration, we study WCPS queries on 2D, 3D, and 4D data sets. Particular emphasis is placed on the computational load queries generate in such on-demand processing and filtering. We look at different techniques and their impact on performance, such as adaptive storage partitioning, query rewriting, and just-in-time compilation. Results show that there is significant potential for effective server-side optimization once a query language is sufficiently high-level and declarative.