Geophysical Research Abstracts,
Vol. 10, EGU2008-A-10955, 2008
SRef-ID: 1607-7962/gra/EGU2008-A-10955
EGU General Assembly 2008

An Efficient Implementation Platform for Coverage Service Standards

P. Baumann

Jacobs University Bremen [p.baumann@jacobs-university.de]

In the modular geo service specification set of the Open GeoSpatial Consortium (OGC, www.opengeospatial.org) the foundation for coverage handling is laid down in the Web Coverage Service (WCS) standard, its current version being 1.1.2. WCS focuses on spatio-temporal retrieval services based on spatial and temporal subsetting, range (i.e., band) subsetting, scaling, reprojection (which may involve resampling and interpolation), and data format encoding for shipping the result.

Actually, in response to many requests WCS is becoming a suite of standards where, based on a common coverage model, server implementors may choose to provide add-on "extension" functionality beyond the "core" functionality every implementation must offer.

One such extension is the Web Coverage Processing Service (WCPS). This recently completed specification offers an XQuery-like coverage request language to send coverage processing requests of unlimited complexity to the server and get back one or more result coverages. Among the examples which the WCPS expressive power allows are WCS functions, image overlaying, search across a coverage set for user-defined properties, statistics, and image/signal processing (such as filter and convolution kernels). Still, the language is "safe in evaluation" - no single request can ever entirely block the server.

The WCPS reference implementation is based on the multi-dimensional raster DBMS, rasdaman. This system allows to store n-D coverages of unlimited size in standard relational databases, and adds an SQL-like raster query language for retrieving and manipulating them. Extensive server-side query optimization is performed. During

optimization, which takes less than a ms on a standard PC, algebraically defined rules are searched to find out possibly more efficient, but semantically equivalent formulations of the query. Raster objects internally are partitioned into tiles; based on a number of tiling strategies tile size and shape can be chosen individually, thereby obtaining a tuning parameter. Raster query processing is based on tile streaming. Frequently, therefore, only one tile at a time needs to be in server main memory. Further, rasdaman makes extensive use of hardware and software parallelization, adapting concepts from GIS, supercomputing, and business data analysis. For extremely large data sets, spatially optimized tape storage management is available.

Benchmarks of rasdaman have shown very positive performance results, and in particular an orders-of-magnitude speedup potential for query optimization.

In our presentation we want to argue and practically demonstrate how databases allows for efficient implementation of value-added coverage services as defined by WCS and WCPS, based on the example of rasdaman which is in operational use since several years and serving, e.g., dozen-Terabytes airborne images, 3-D geophysics data, and 4-D climate simulation results. As a proof of concept, a Web Map Service (WMS) has been implemented on top of WCPS. An on-line demo is under progress at www.earthlook.org.

The author is co-chair of the WCS and Coverages Working Group, chair of the WCPS group, and principal architect of rasdaman.