Title: "Using spatial databases to efficiently manage 3D objects at a national level"

By Hans Viehmann, Oracle

Abstract: "Management of the various kinds of 3-dimensional geospatial data and the associated metadata at a national level almost inevitably poses a significant challenge due to the large data volumes involved. Particularly the collection of LiDAR data at high resolution has led to a massive growth in the size of datasets, but also other 3D data types such as 3D vector data used in detailed 3D city models are contributing significantly to the growing data volumes, particularly when used in conjunction with geo-referenced raster imagery for terrain or facade rendering.

These large datasets are usually visualized or processed as a combination of 3D data, various kinds of 2D data or enhanced through attribute information. Both visualization and processing benefit from having fast and simple access to the different kinds of data, independent of data volume. This is a major strength of today's spatial databases.

In this paper we are describing the use of Oracle Spatial and Graph to manage huge volumes of 3D geospatial data. We have been running a series of benchmarks jointly with Rijkswaterstraat of the Netherlands, the Netherlands eScience Center, Fugro GeoServices and TU Delft using the second national height map of the Netherlands (AHN2) as a basis. The dataset consists of some 640 billion measurement points and was originally delivered in over 60000 LAS-formatted files. We investigated the use of compression techniques in the database, data partitioning to allow for spatial clustering and multi-processor support for scalability. For this purpose we loaded the dataset onto an Oracle Exadata Database Machine and executed a series of queries against the system.

The results indicate that this combination of hardware and software is capable of delivering fast response times for data retrieval even for very large data volumes. At the same time it allows the execution of procedural logic inside the database without transferring massive amounts of data across the network for analysis. This is beneficial for the processing of point cloud data (point classification, normal vectors) and is likely to become even more important once the datasets are partly updated and change detection is required. We will discuss the various aspects of data management and processing and provide an outlook on possibilities and future work."