

Using a Space Filling Curve for the management of dynamic point cloud data in a Relational DBMS

Stella Psomadaki
TU Delft

Abstract:

The rapid developments in the field of point cloud acquisition technologies have allowed point clouds to become an important source of information for many applications. One of the newest applications of point clouds concerns the monitoring of the coast. Many countries, among which the Netherlands, use this source of data in order to determine the changes in coastal elevations. This means that point clouds are collected every hour, day, month, year; ultimately talking about dynamic point clouds. To be able to efficiently use this plethora of data, the management of those point clouds, dynamic or not, is proven to be crucial.

Point clouds, like the majority of geodata, have been traditionally managed using file-based solutions. Nevertheless, the last years database solutions have emerged. Typical examples are the point cloud extensions for PostgreSQL and the Oracle Database. Both options use a similar block-based organisation. In addition to the block based organisations, point clouds can also be managed using a flat table where each point is stored in a separate row. While the first approach is very scalable and efficient, the second is easier to implement and to update. To make the flat model scalable, a Space Filling Curve (SFC) can be used to cluster the data. Nonetheless, both approaches in their current forms, are not suited for the management of dynamic points. The reason for this is the fact that they do not consider the time dimension as part of the organisation and further insertions for the block-based approaches are not straightforward.

Within this presentation a SFC approach for managing dynamic point clouds is investigated. For this, the flat model approach using an Index Organised Table (IOT) within a Relational Database Management System (RDBMS) is used. Two variants coming from two extremes of the space - time continuum are then taken into account. In the first approach, space and time are both used within the SFC (integrated approach), while in the second one, time dominates over space (non-integrated approach). Along these two approaches, two treatments of the z dimension are, also, studied: as attribute or as part of the SFC. In addition to that, building on the coastal monitoring applications, the most important queries are identified: space - time, only time, only space.

The efficiency of the implemented methodology is tested through the execution of a benchmark. Using two use cases coming from coastal applications, the benchmark is executed once for daily and once for yearly data. The results show that the SFC approach is an appropriate method for managing dynamic point clouds. Furthermore, the integrated approach is the most suitable way to proceed. Achieving scalability, time efficiency and dynamic insertions can be achieved for various use cases.




Figure 1 The integrated space and time approach (z is not included in the SFC)




Figure 2 The non-integrated space and time approach (z is not included in the SFC)