

Distributed SECONDO: A highly available and scalable system for spatial data processing

Jan Kristof Nidzwetzki and Ralf Hartmut Güting


Faculty of Mathematics and Computer Science, FernUniversität in Hagen, Germany



Distributed Secondo

Distributed SECONDO

The extensible, highly available, scalable DBMS



Motivation

Problem: Database management systems (DBMS) have to deal with rapidly growing amounts of data. Nowadays, distributed key-value stores are often used to store huge amounts of data. The current situation can be roughly summarized as follows.

DBMS ...

- ▶ ... scale poorly.
- ▶ ... provide a lot of functions (e.g. joins or spatial joins) to analyze data.

Key-Value stores ...

- ▶ ... scale well.
- ▶ ... provide very limited functions regarding data analyzes. Functions like joins have to be implemented repeatedly in the application code.

Solution: Couple a DBMS with a key-value store to create a highly available and scalable data base management system.

System overview

SECONDO is an extensible DBMS developed at FernUniversität Hagen. The system is designed with a focus on supporting spatial and spatio-temporal data.

DISTRIBUTED SECONDO is a distributed version of SECONDO. The system couples SECONDO with the key-value store Apache Cassandra. All the details about data distribution, parallel query processing and fault tolerance are encapsulated in one software component.

As a result, almost all of the data models and functions provided by SECONDO can be used in a distributed and scalable manner without changing their implementation.

The system consists of three node types:

- ▶ Storage Nodes (SNs) – Responsible for storing data.
- ▶ Query Processing Nodes (QPNs) – Responsible for query processing.
- ▶ Management Nodes (MNs) – Responsible for importing and exporting data. Also, the query plans for the QPNs can be managed.

DISTRIBUTED SECONDO is scalable in two dimensions: (i) By adding storage nodes, bigger amounts of data can be stored. (ii) By adding query processing nodes, data can be analyzed faster.

Data flow

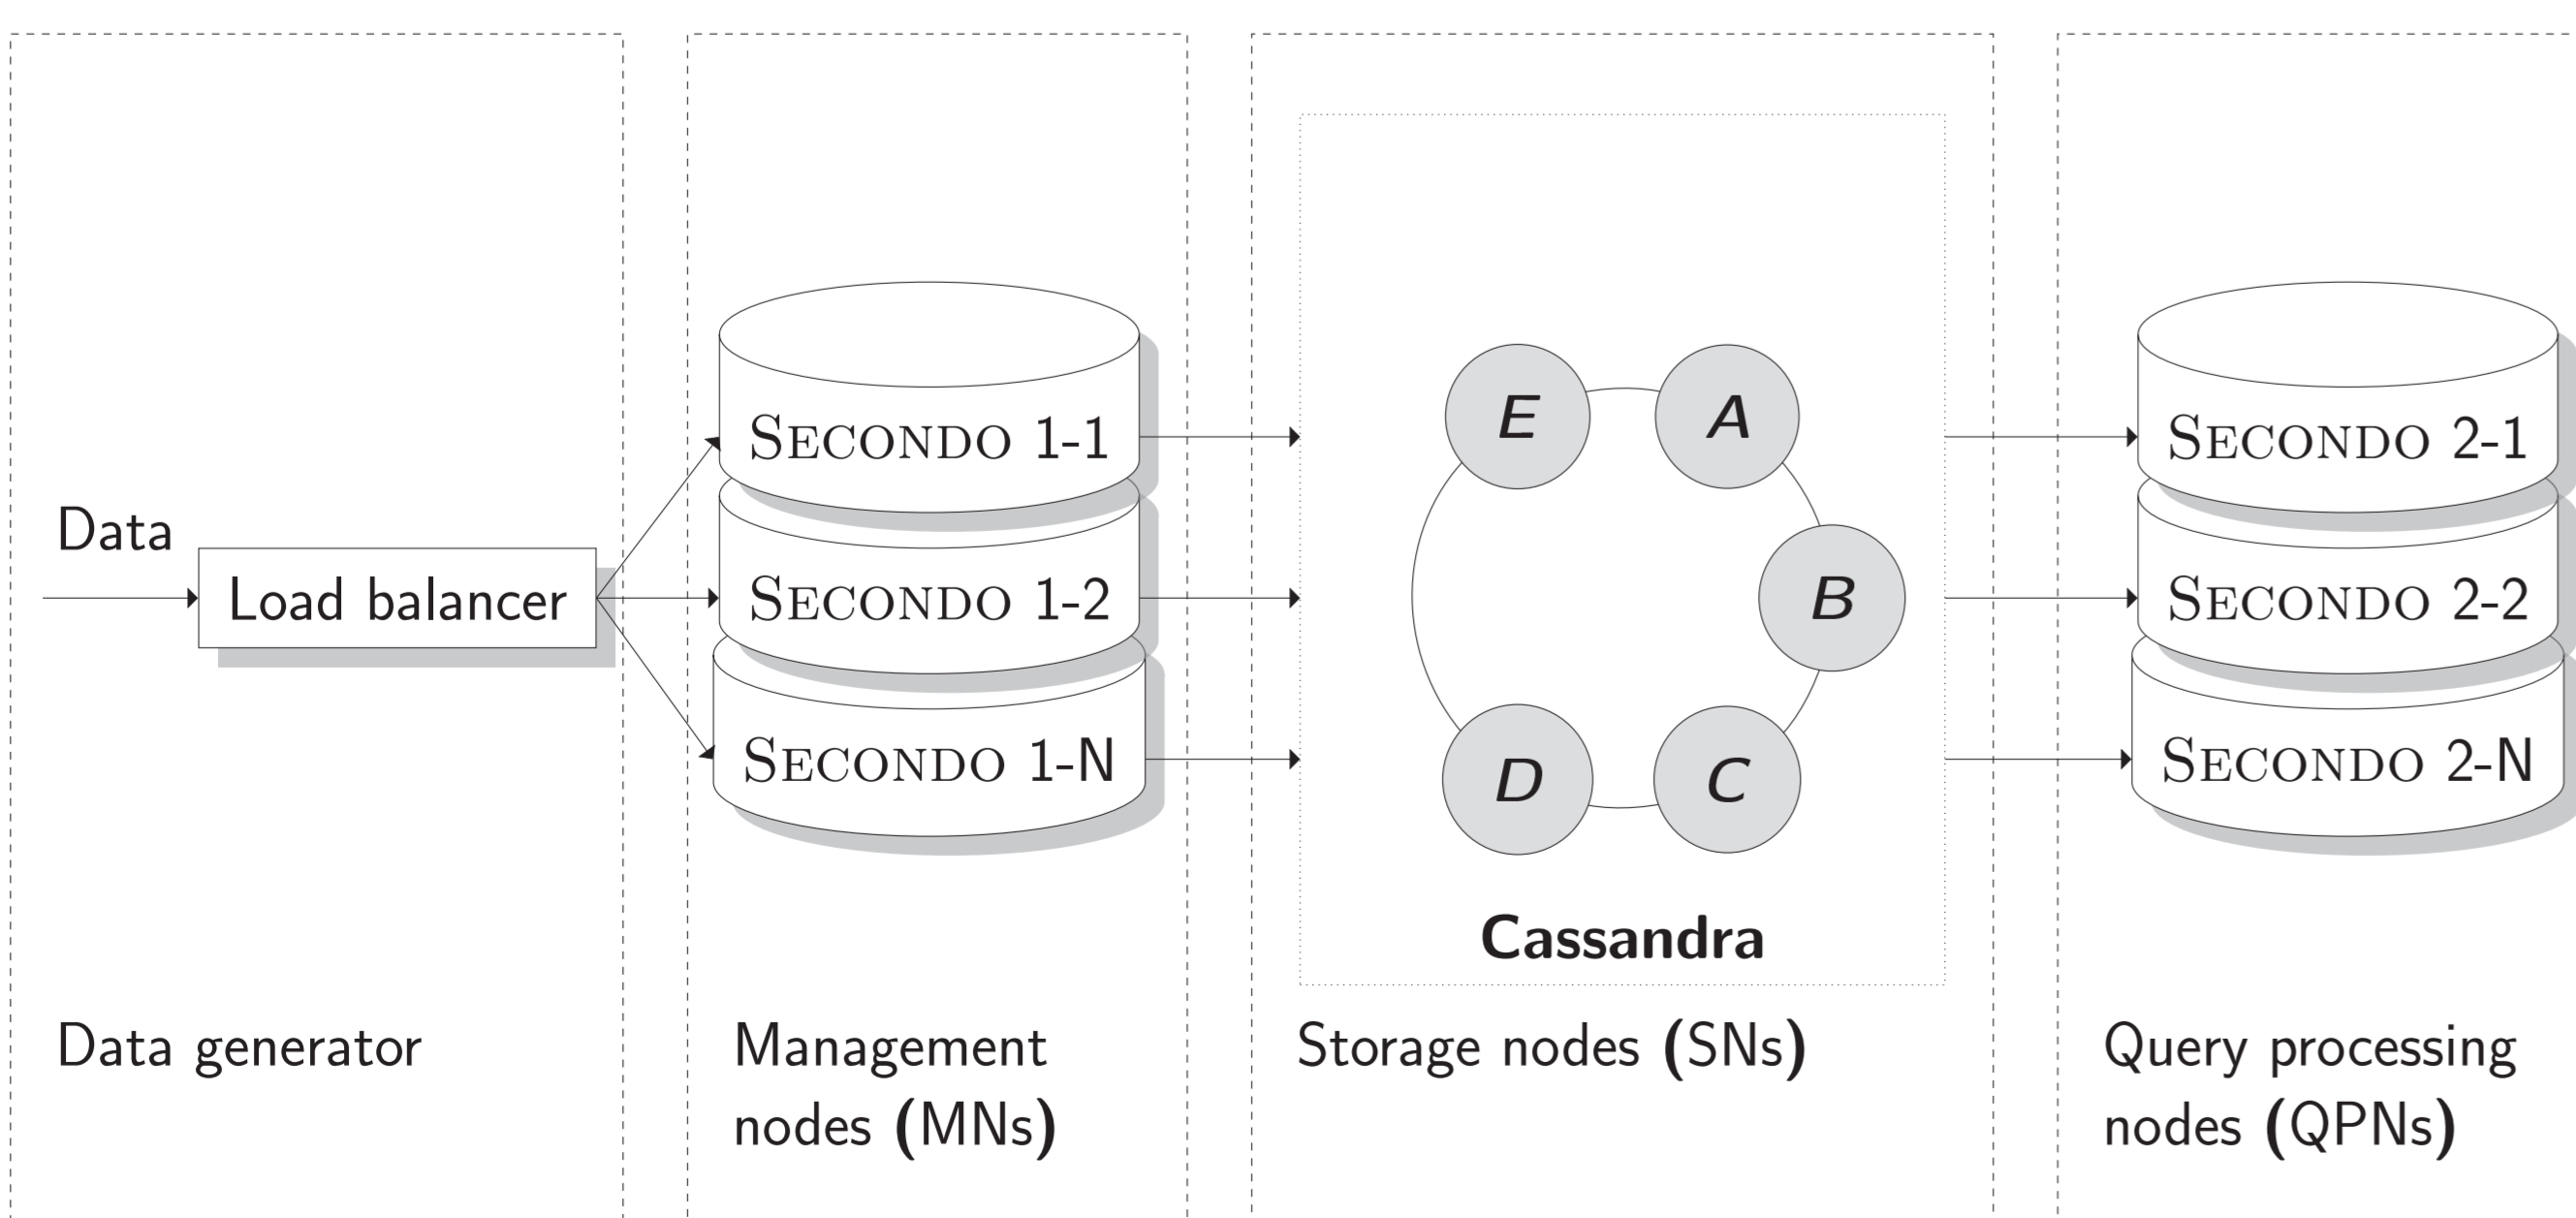


Figure: An overview of the data flow in DISTRIBUTED SECONDO. New data is processed by a load balancer. The load balancer distributes the data on the MNs. These MNs process the data and store them on the SNs. In a further step, the data is analyzed on the QPNs.

BerlinMODPlayer

BerlinMOD is a benchmark for spatio-temporal database management systems. The benchmark contains a data generator which creates trips of moving vehicles within Berlin. BerlinMODPlayer is a player for such data. The player generates a real time stream of GPS coordinates.

In a real world scenario, such data is produced by a fleet of moving vehicles. Each one is equipped with a GPS receiver and sends out coordinate updates to a central system.

Position updates are represented as lines. One position update has the following format:

Example: Time stamp \downarrow $28-05-2007\ 10:00:14$, Moid \uparrow 10 , Tripid \downarrow 1313 , X coordinate \uparrow 13.2967 , Y coordinate \downarrow 52.4502

Demo I: Scalable processing of GPS updates

DISTRIBUTED SECONDO is used in this demonstration to process a GPS coordinate stream, generated by BerlinMODPlayer. The coordinate stream is read by a MN and converted into tuples. After conversion, the MN stores the tuples onto the SNs.

Demo II: Performing a distributed spatial join

In this demonstration, a spatial join is performed. For this purpose, two relations with spatial data are stored on the SNs. The relations contain the roads and the forests of an area in Germany. The spatial join calculates which roads lead through a forest.


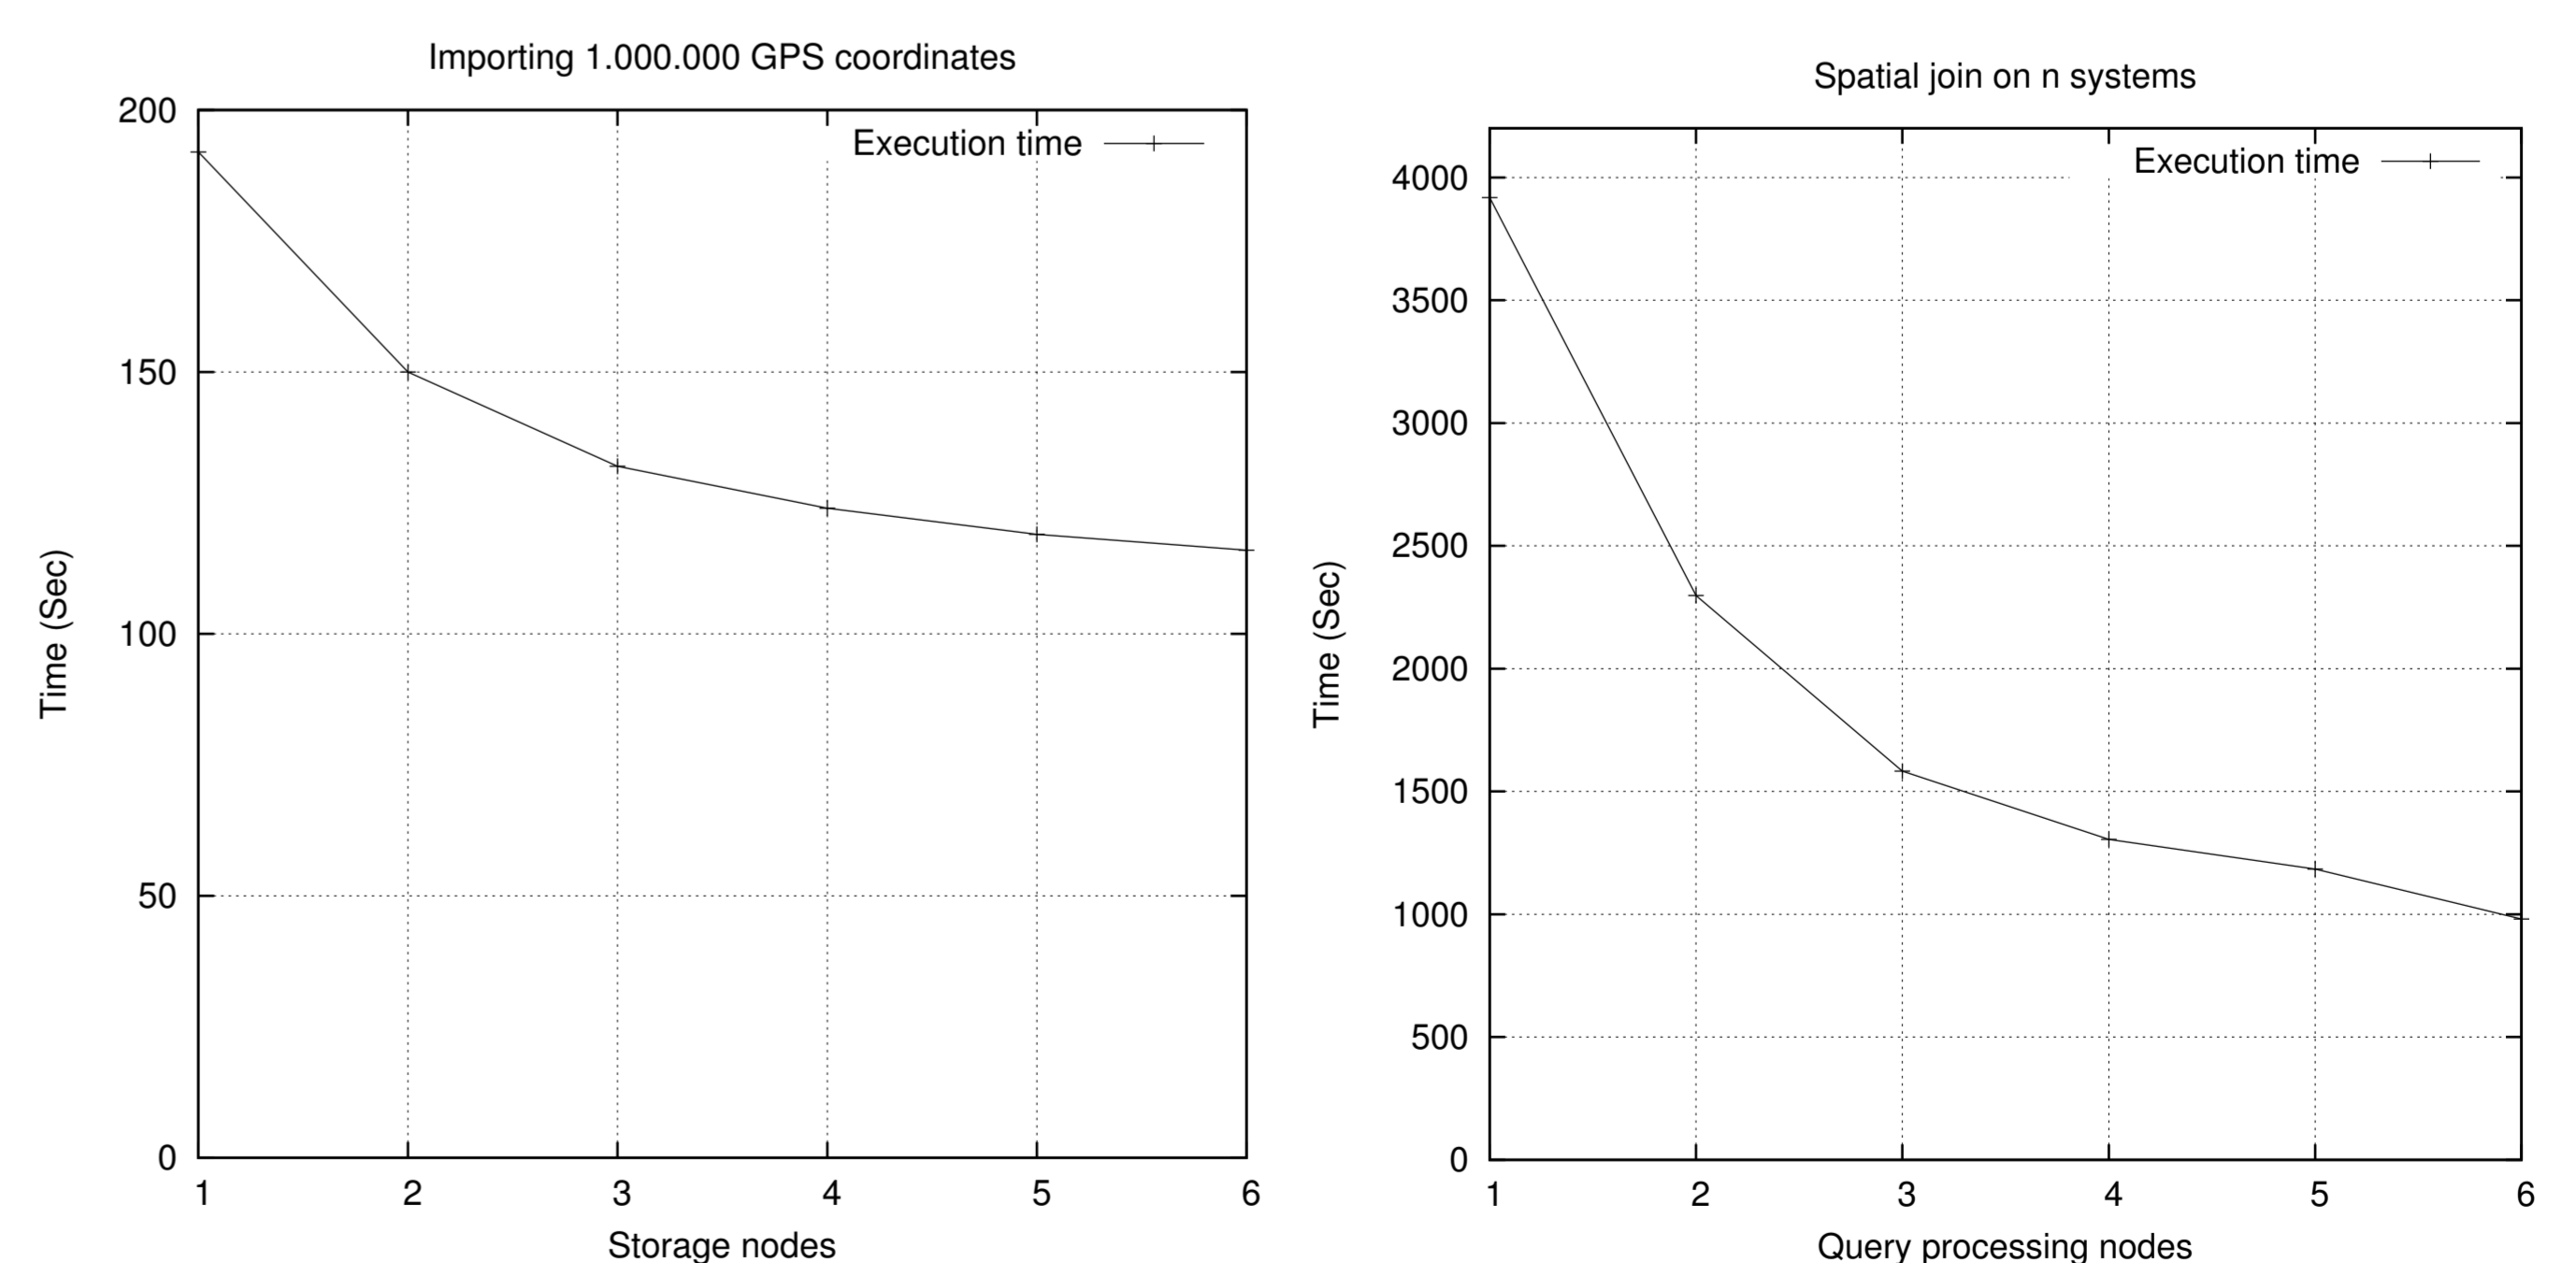


Figure: Two relations with spatial data stored in Cassandra. The spatial data is overlapped with a grid. The content of corresponding cells in both relations are stored at the same position in Cassandra.

Evaluation



Conclusion

DISTRIBUTED SECONDO ...

- ▶ ... is a highly available, distributed and scalable DBMS.
- ▶ ... can be used to execute spatial joins.
- ▶ ... can be scaled up and down by adding or removing SNs and QPNs.
- ▶ ... can also be used to process other kinds of data (e.g. relational).