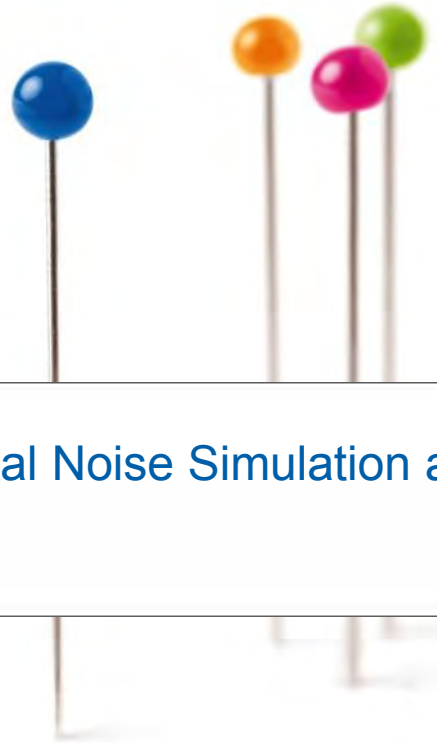

Creating a 3D Model for Environmental Noise Simulation at the German Federal Railways Agency

Creating a 3D Model for Noise Simulation at the German Federal Railways Agency

OVERVIEW

- disy is a leading German software and consulting company
- specialized in management and analysis of environmental and spatial data
- key customers are big public authorities
- Foundation: 1997, Employees: 55, Location: Karlsruhe, Germany

CHALLENGES / OPPORTUNITIES

- Large amount of complex, heterogeneous, redundant and incorrect data has to be implemented in a short time period
- Traceability, documentation and repeatability of the process
- Problems appear after the project has already started

SOLUTIONS

- Oracle Database 11g and 12c Enterprise Edition
- Spatial Option with Geocoder, Spatial Vector Acceleration, …
- Java with open source and own GIS libraries
- Solaris

RESULTS

- Despite delays on the customer side and a very ambitious schedule the project was finished in time
- Previously unknown data problems solved due to adoptable procedures
- Reusable process for regularly recurring tasks
- Smart combination of technologies Solaris, Java, Oracle Spatial
- No need of a traditional GIS in core processes

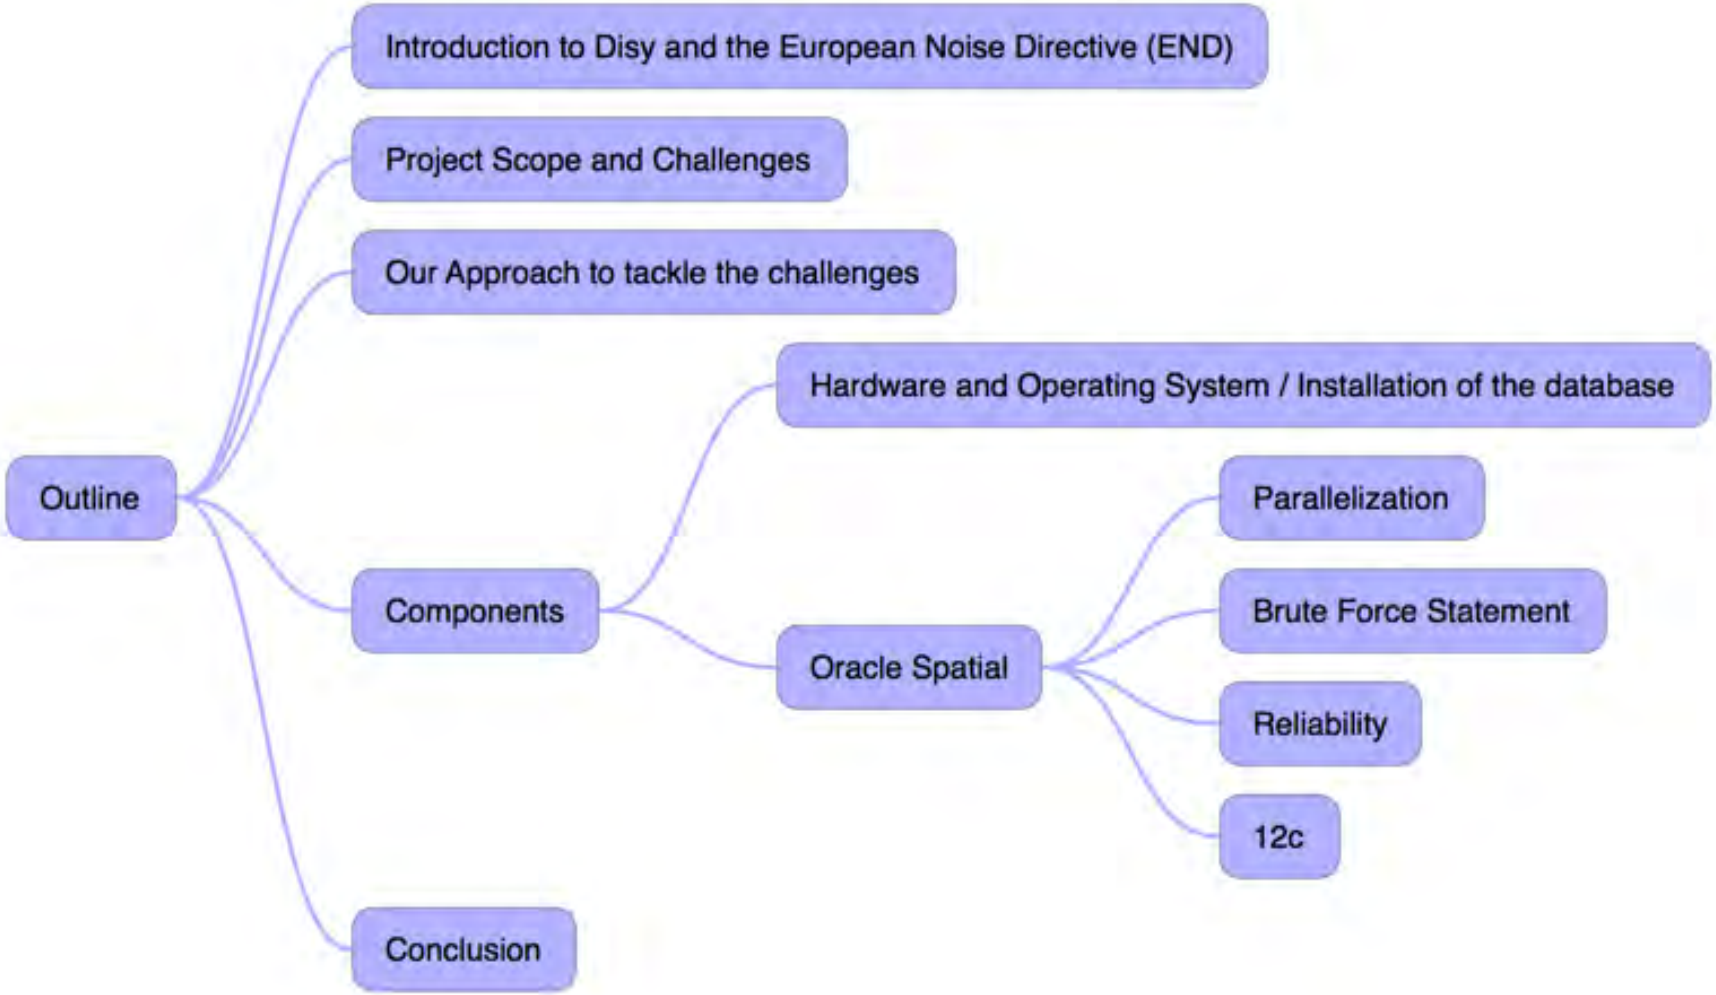
Outline
Introduction to Disy and the European Noise Directive (END)
Project Scope and Challenges
Our Approach to tackle the challenges
Components
Hardware and Operating System / Installation of the database
Oracle Spatial
Parallelization
Brute Force Statement
Reliability
12c
Conclusion

disy is…

… a leading German software and consulting company for **managing and analysis of environmental and spatial data**

Who we are

- Foundation: 1997, Employees: 55, Location: Karlsruhe, Germany
- Our key customers are public authorities for environment, agriculture, transportation or infrastructure organizations like railroad companies
- In many Projects we implement large scale information systems based on our core technology to help our clients in a wide variety of reporting processes within the context of the European Framework Directives on Water, Flood Risk, Noise Pollution, Agricultural Promotion

The Environmental Noise Directive (2002/49/EC)

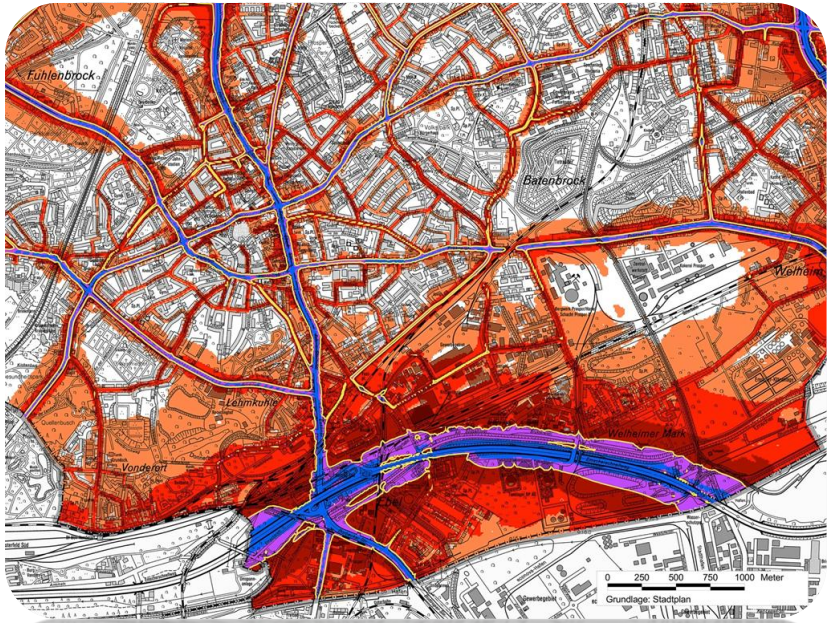

Directive of European Commission

Goals of the directive:

- Monitoring the environmental problem of noise pollution
- Informing and consulting with the public
- Requiring action plans to address local noise issues
- Developing a long-term EU strategy

How to transform raw data to noise maps?

disy

How to transform raw data to noise maps?

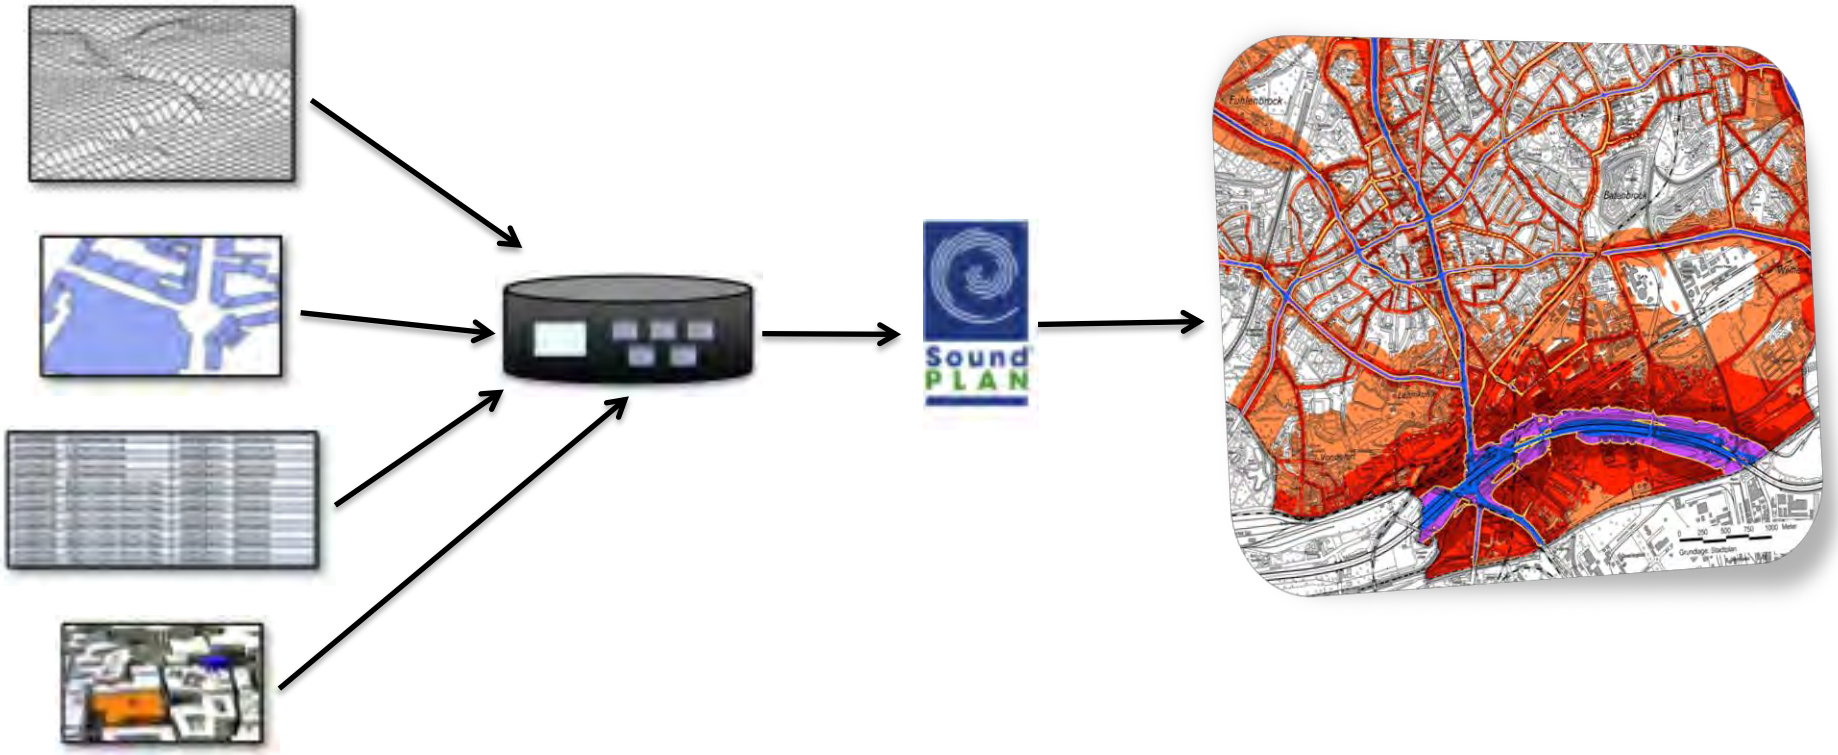

disy

Project Scope – Example Processing

- Correction of geometry for SDO-conformity
- Data reduction through simplification
- Transformation to a single coordinate system
- Verification for completeness
- Union of single data sets
- Geocoding (usage of building)
- Triangulation and thinning of DTM
- Building level estimation
- Data reduction to mapping area

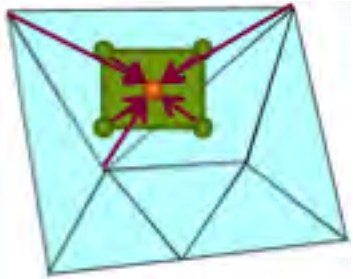

Challenges in the project (1)

- about 150 GB input data
- 30.000 single files with about 90 million buildings
- More than 30 million addresses
- More than 75 GB DTM (10 m)

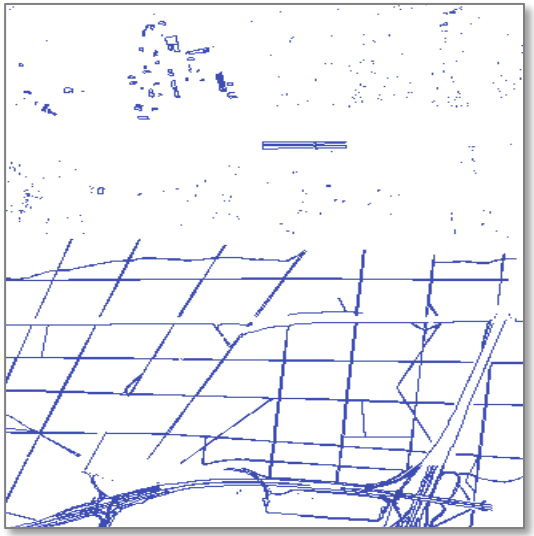

- More than 50 different sources
- different data formats
- redundancies, inconsistencies, contradictions,
- unknowns, below standard quality

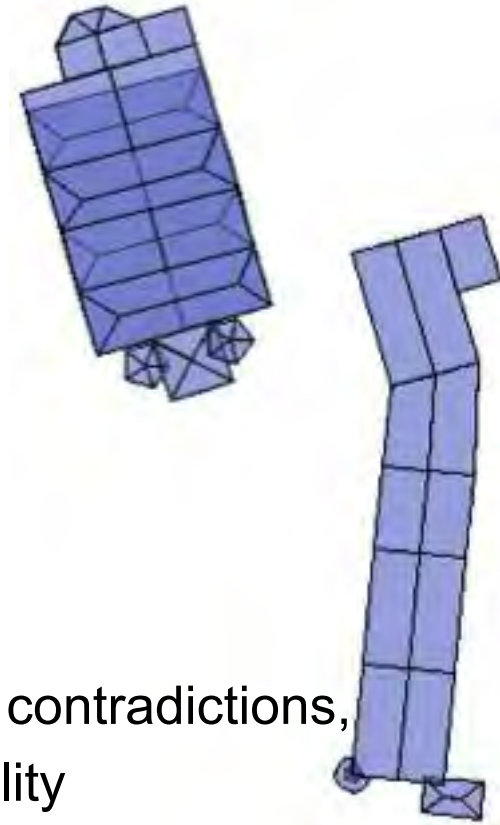

Challenges in the project (2)

Complexity

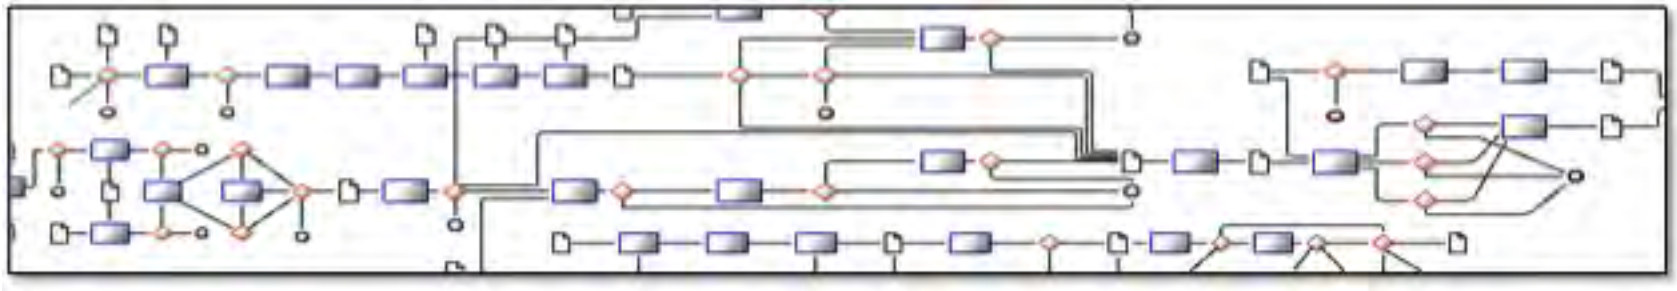

> 250 process steps

Timeframe

- About 6 Months
- Work sample after 3 Months

disy

Our approach

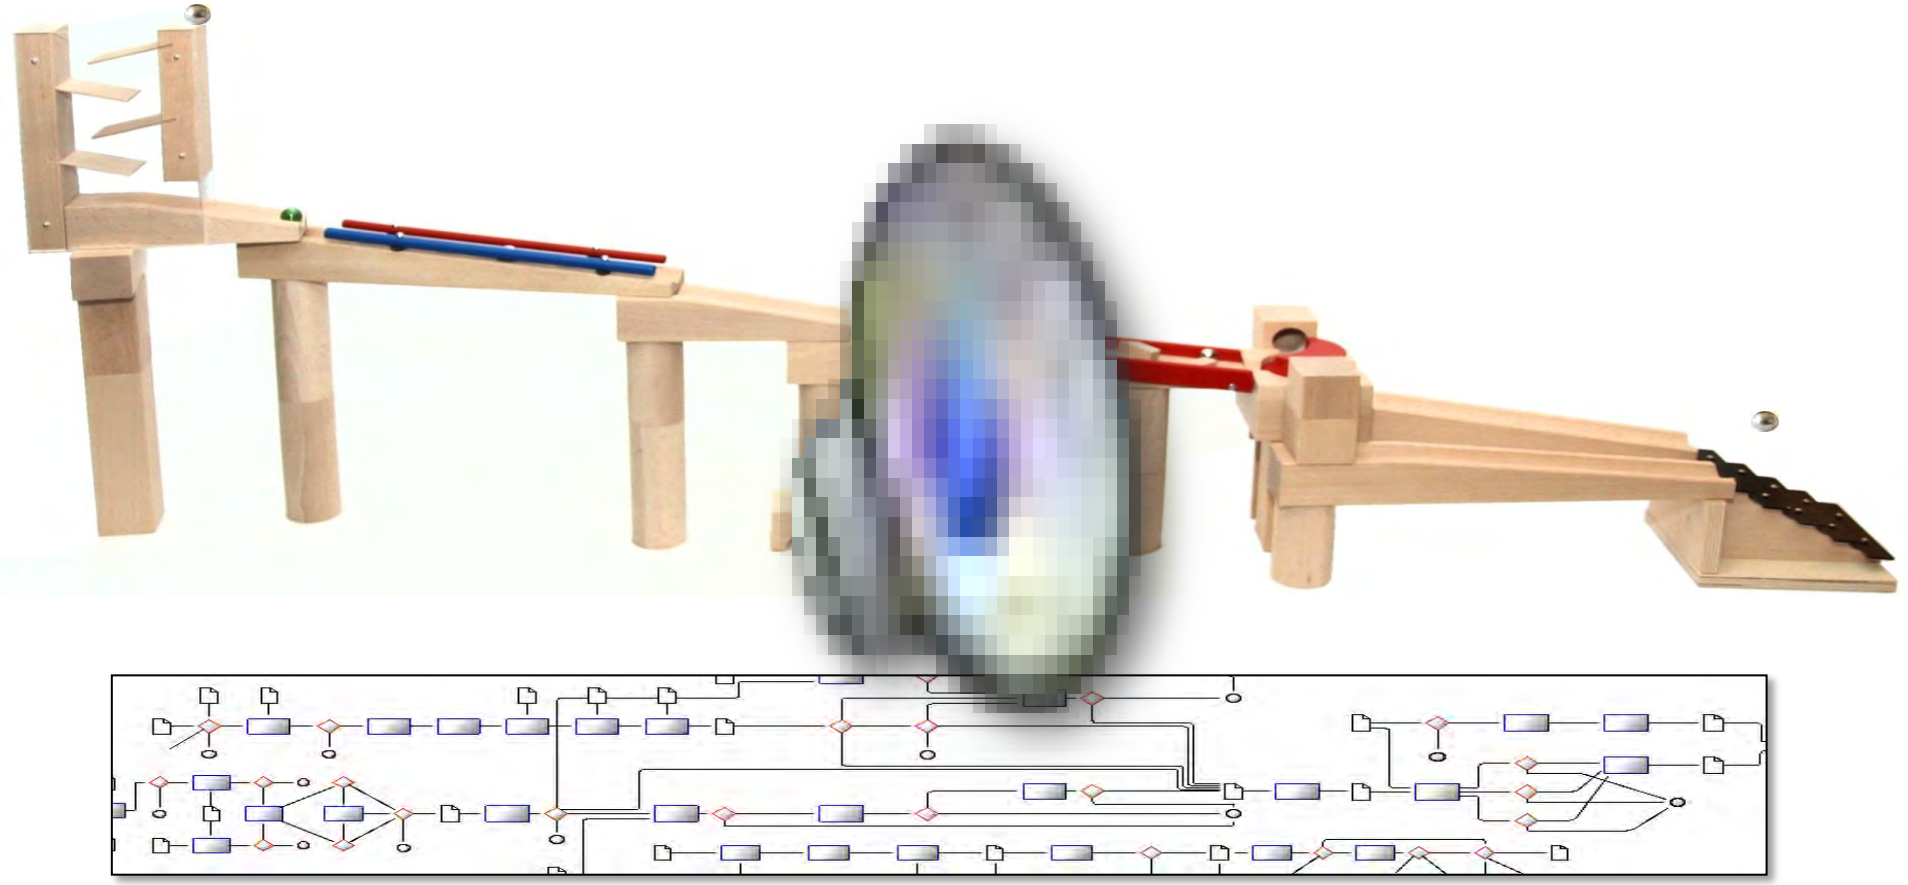

Our rationale

All Data comes into and stays in the database

ORACLE®

Building blocks:

PL/SQL: 300+ algorithms

Java: GISterm, GeoTools, …

…others…

disy

Rolling through the marble run

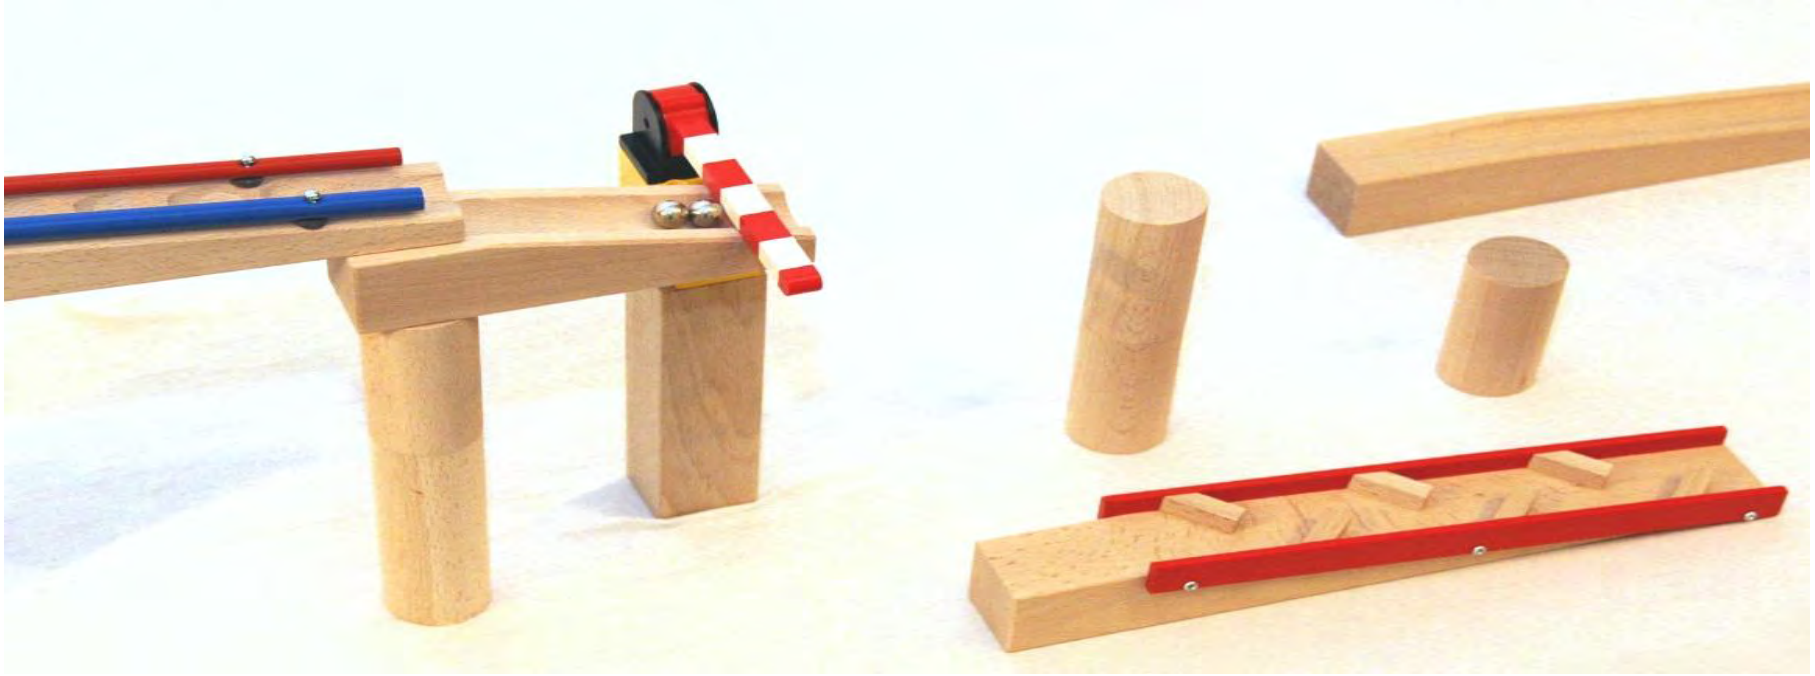

Data needed to be processed while marble run was being built

Customers decisions led to algorithm changes ➜ back to start

Our approach – technical overview

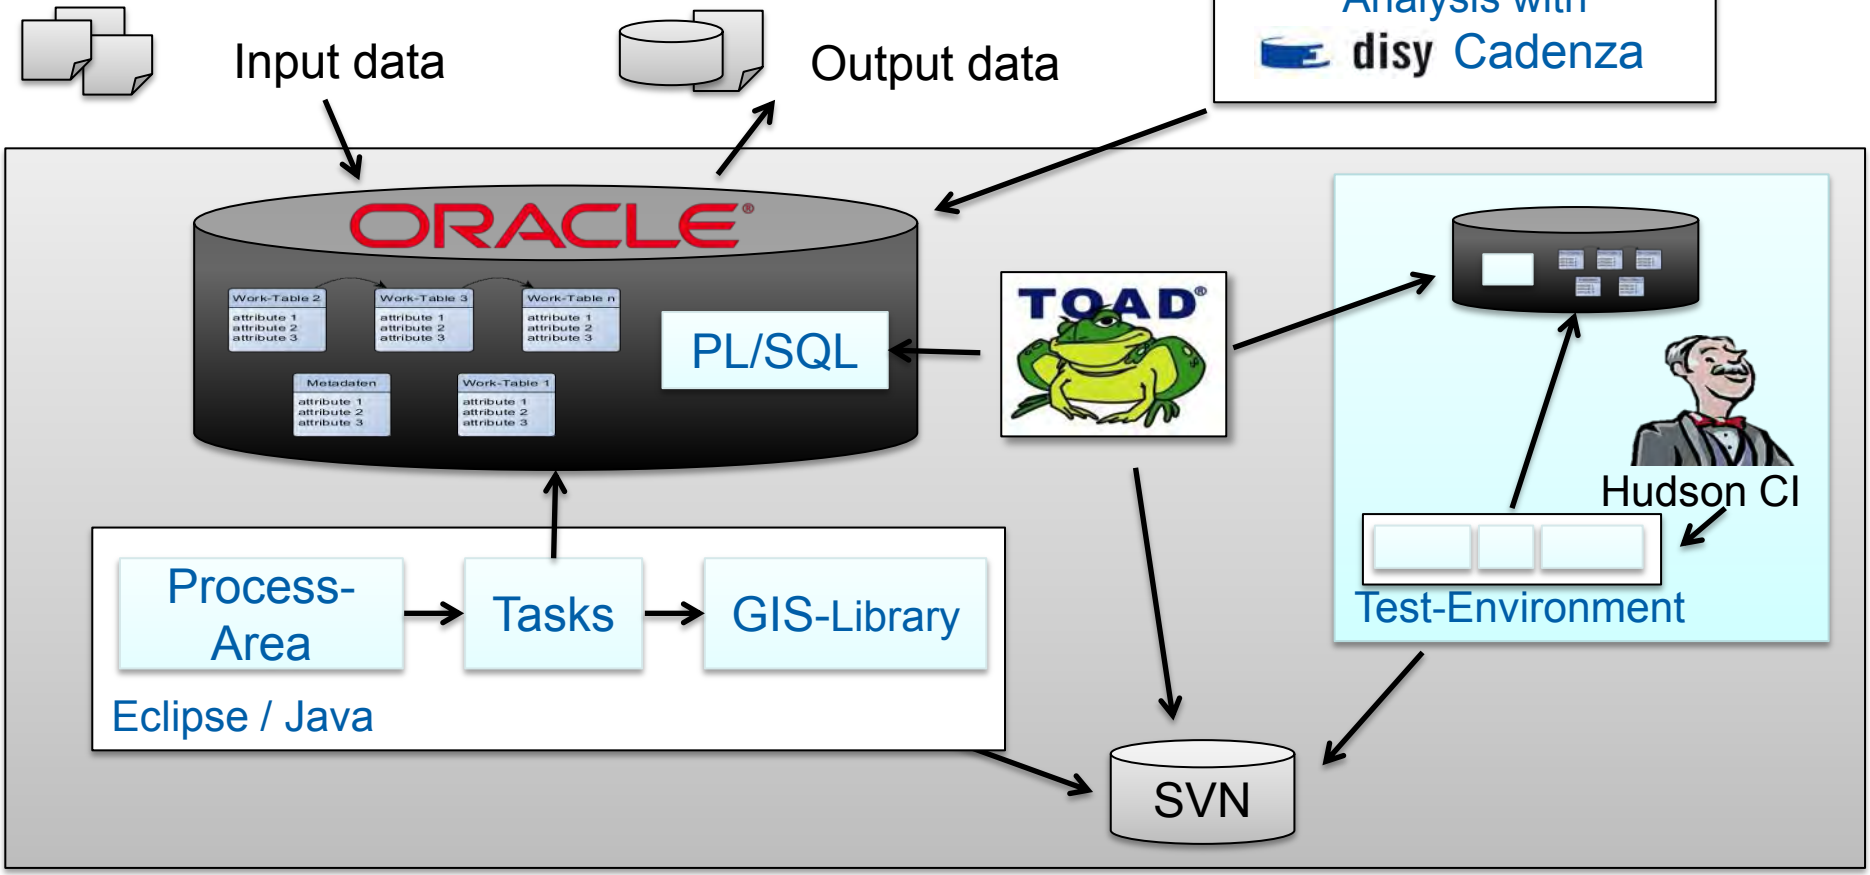

Process-Steps – Condition of data set as status table in DB

Status Flag

Process steps: Timestamp

IDENTIFIER	LAYER_ID	STATE_ID	CREATE_DATE	LAYER_CREATE_DATE	REGISTRATION_DATE
164	134503	-1	30.06.2011 08:29:48	15.08.2011 23:55:25	16.08.2011 06:53:37
169	128940	2	30.06.2011 08:29:49	15.08.2011 23:55:02	16.08.2011 06:53:37
193	120399	2	30.06.2011 08:29:53	15.08.2011 23:55:01	16.08.2011 06:53:37
381	86478	-1	14.07.2011 11:25:43	15.08.2011 23:55:05	16.08.2011 06:53:37
382	86130	-1	14.07.2011 11:25:43	15.08.2011 23:55:17	16.08.2011 06:53:38
383	52261	-1	14.07.2011 11:25:43	15.08.2011 23:55:16	16.08.2011 06:53:39
386	84116	2	14.07.2011 11:25:44	15.08.2011 23:55:49	16.08.2011 06:53:41
393	51388	-1	14.07.2011 11:25:44	15.08.2011 23:56:38	16.08.2011 06:53:41
418	96536	-1	14.07.2011 11:25:45		
420	112277	-1	14.07.2011 11:25:45	16.08.2011 00:21:35	16.08.2011 06:54:45
424	134176	-1	14.07.2011 11:25:46		
426	134277	-1	14.07.2011 11:25:46		
428	120461	-1	14.07.2011 11:25:46		
430	148611	-1	14.07.2011 11:25:46	16.08.2011 00:25:07	16.08.2011 06:54:59

IDENTIFIER	NAME
-1	UNKNOWN
1	NEW
3	OLD
4	IGNORE
2	CURRENT

Task = Java-Class

- Executes a Process-Step
- Processes a single, atomar data entity
- Uses PL/SQL, Java GIS-Routines or Java+SQL, …

Task-Execution: simple, Java-based Execution-Environment

disy

Hardware and Operating System

Bottleneck: IOPS

Solution: Whole Database works from RAM Disk

Prioritize performance vs. security

Dedicated Database Server:

- 4 x Intel(R) Xeon(R) CPU E5
- 768GB RAM
- Solaris
- ZFS

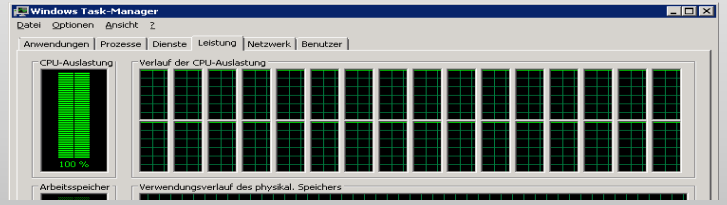

disy

Parallelization

Goal:

- Optimal CPU Usage
- No delays because of unfinished process steps

Strategy:

- All tasks may run parallel
- Independent data is processed parallel

Implementation:

- Break up in small, independent steps
- Status: “Paused“
- Best result: 24 Threads with 32 CPUs

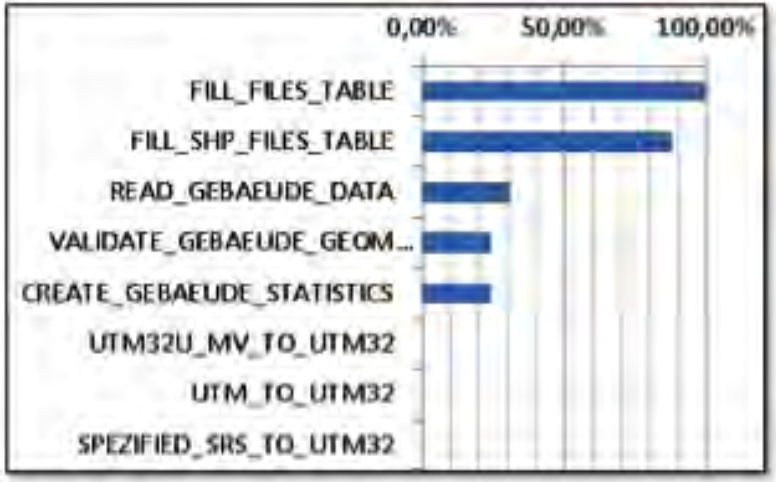

disy

Oracle Spatial Geocoder: Assignment of addresses to buildings

Disadvantages: High initial efforts

- Documentation: Focus on using an existing Geocoder dataset
- Did not work out of the box (Patch p15926096)
- Setting up and populating the complex schema is not trivial

Advantages:

- Fuzzy matching with supporting “Match vector“
- Allows detailed interpretation of result quality
- Usable with incomplete data

→ ***Proved to be a good solution for geocoding our data***

2110766825	21	Hasensprung	10318	Berlin	Karlshorst
2110766824	17	Hasensprung	10318	Berlin	Karlshorst
2110766823	11	Hasensprung	10318	Berlin	Karlshorst
2110766822	9	Hasensprung	10318	Berlin	Karlshorst
2110766819	1	Hasensprung	10318	Berlin	Karlshorst
2110863678	40	Kol.Frieden Neuer Weg	12099	Berlin	Westend
2110863677	38	Kol.Frieden Neuer Weg	12099	Berlin	Westend
2110863676	31	Kol.Frieden Neuer Weg	12099	Berlin	Westend
2110863675	13	Kol.Frieden Neuer Weg	12099	Berlin	Westend
2110863674	12	Kol.Frieden Neuer Weg	12099	Berlin	Westend
2110863673	11	Kol.Frieden Neuer Weg	12099	Berlin	Westend
2110850781	33	Kol.Frieden Neuer Weg	12099	Berlin	Westend
2110850773	25	Kol.Frieden Neuer Weg	12099	Berlin	Westend
2110850758	10	Kol.Frieden Neuer Weg	12099	Berlin	Westend
2110850757	9	Kol.Frieden Neuer Weg	12099	Berlin	Westend
2110850756	8	Kol.Frieden Neuer Weg	12099	Berlin	Westend
2110826078	21	Grabbeallee	13156	Berlin	Niederschönhausen
2110826079	27	Grabbeallee	13156	Berlin	Niederschönhausen
2110826081	30	Grabbeallee	13156	Berlin	Niederschönhausen
2110826085	43	Grabbeallee	13156	Berlin	Niederschönhausen
2110826086	66	Grabbeallee	13156	Berlin	Niederschönhausen
2110826087	80	Grabbeallee	13156	Berlin	Niederschönhausen
2110861414	53	Grabbeallee	13156	Berlin	Niederschönhausen
2110850548	39	Kol. Eichtal	14050	Berlin	Charlottenburg
2110850547	38	Kol. Eichtal	14050	Berlin	Charlottenburg
2110850546	36	Kol. Eichtal	14050	Berlin	Charlottenburg
2110850545	35	Kol. Eichtal	14050	Berlin	Charlottenburg
2110850544	34	Kol. Eichtal	14050	Berlin	Charlottenburg
2110850543	31	Kol. Eichtal	14050	Berlin	Charlottenburg
2110850542	30	Kol. Eichtal	14050	Berlin	Charlottenburg

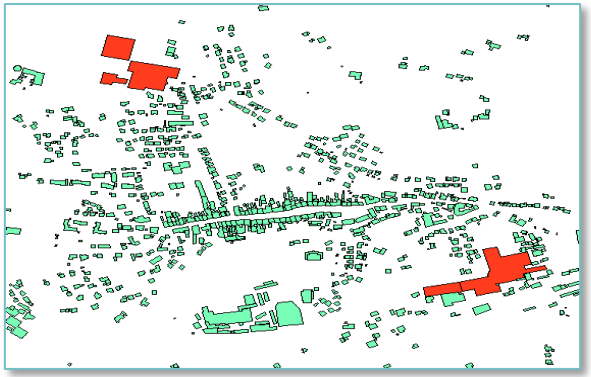

Reliability of Oracle Spatial

Bugs in Oracle Spatial:

- Software projects of this size inevitably lead to bugs!

Positive experiences:

- Most bugs have not been critical
- Most bugs have been resolved in advance, patches were available
- There are often workarounds (no official documentation)
- When in doubt, use “Hints“

Strategy for the project:

- Don‘t waste too much time on bugs!
- Instead use workarounds or other technologies

→ ***Considering the number of features provided, Oracle Spatial proved to be a sufficiently reliable solution***

Performance of Oracle Spatial – process steps (example)

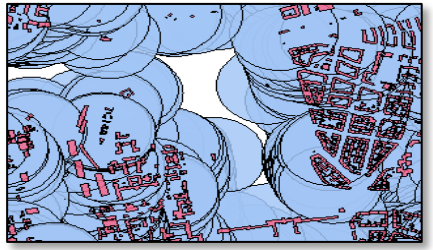

Merging
of buffered building
geometries

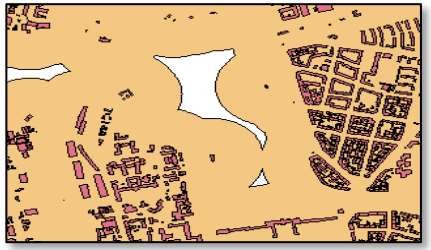

```
PROCEDURE merge_big(v_table_id IN NUMBER)
IS
  v_tablename VARCHAR2(38) := 'tmp_simple_buffered_' || v_table_id;
  l_merged_buffer MDSYS.sdo_geometry;
BEGIN
  EXECUTE IMMEDIATE
    'select SDO_AGGR_UNION(MDSYS.SDOAGGRTYPE(buffered, 0.005)) GEOM
   FROM ' || v_tablename
    INTO l_merged_buffer;
  EXECUTE IMMEDIATE 'delete from tmp_big_merged_buffers where identifier = :table_id'
    USING v_table_id;
  EXECUTE IMMEDIATE
    'insert into tmp_big_merged_buffers(identifier, geometry) values (:identifier, :geometry)'
    USING v_table_id, l_merged_buffer;
END;
```

Example:
365.441 building geometries (3,3 million anchor points)
Duration: 3,5 Days

→ ***Very high performance in complex scenarios***

disy

Some words about 12c …

- Faster (sorry – no measurements)
- No show stopper (in spite of it's a .1 release)
- Bugs in Buffer and aggr_union (should be fixed already)
- We did not use Spatial Vector Acceleration
- Some nice addition to PL/SQL

- TIN and Point Clouds
  - interesting for terrain model, laser scanning and GPS points
  - Needs more experimentation on our site before using it in such a demanding project

Project result

Project closure: in time and in budget

→ Highly automated data processing can be an alternative to traditional approaches for data processing

→ Oracle Spatial is a powerful and reliable core component of such a solution

Thank you for your attention

Your contact at Disy:

Marcus Briesen

Dipl.-Informatiker

Leiter Beratung und Entwicklung

Mitglied der Geschäftsführung

Disy Informationssysteme GmbH

Ludwig-Erhard-Allee 6
76131 Karlsruhe, Germany
www.disy.net

Tel. +49 721 16006-221
Fax +49 721 16006-05
marcus.briesen@disy.net

Abstract

- Like many public bodies and private entities in Europe, the German Federal Railways Agency has to provide a noise emission map in order to comply with the EU Environmental Noise Directive (END). This directive is part of a long-term EU strategy to reduce the exposure of EU citizens to noise.
- For this purpose, a complex 3D model had to be created which was then used to simulate noise distribution along railway lines. The underlying datasets were delivered by various entities from the different states (Länder). In total, datasets from over 50 sources with inconsistencies, redundancies and errors in all kinds of formats had to be merged.
- A “data processing factory” was implemented to conflate and cleanse over 150 GB of geospatial data, including some 30,000 files with 90 million buildings, over 30 million addresses, and a terrain model dataset of more than 75GB in size. The data processing factory is based on a Java framework that controls all processing flows, in conjunction with other tools such as Hudson and a comprehensive PL/SQL library in the database. This library is built on Oracle Spatial and Graph with Oracle Database 12c, using various advanced capabilities such as the geocoding engine.
- In this session we will describe the data processing factory in more detail and share our experience from its implementation and use. We will showcase the possibilities and limits of Oracle’s spatial technologies in this complex real world scenario, and provide insight into the new Oracle Spatial and Graph features in version 12c.

Please list 3 learning objectives for this session

1. Understand why complex spatial processing should be executed inside a Geodatabase like Oracle Spatial

2. Learn about multiple geoprocessing steps that can be combined using a steering framework like the Disy Spatial Workbench
3. Learn about features, stability, performance, and robustness of Oracle Spatial and Graph in the Oracle Database 12c release