Spatial Databases - a look into the future

Mario Miler and Damir Medak

Faculty of Geodesy
University of Zagreb

November 26, 2010

- Introduction
- Spatial databases
  - geography data type
  - curves
  - raster data type
- NoSQL
  - document datastore (CouchDB, MongoDB)
  - graph datastore (Neo4j)

- development of spatial databases during the last few years has been rapidly increasing

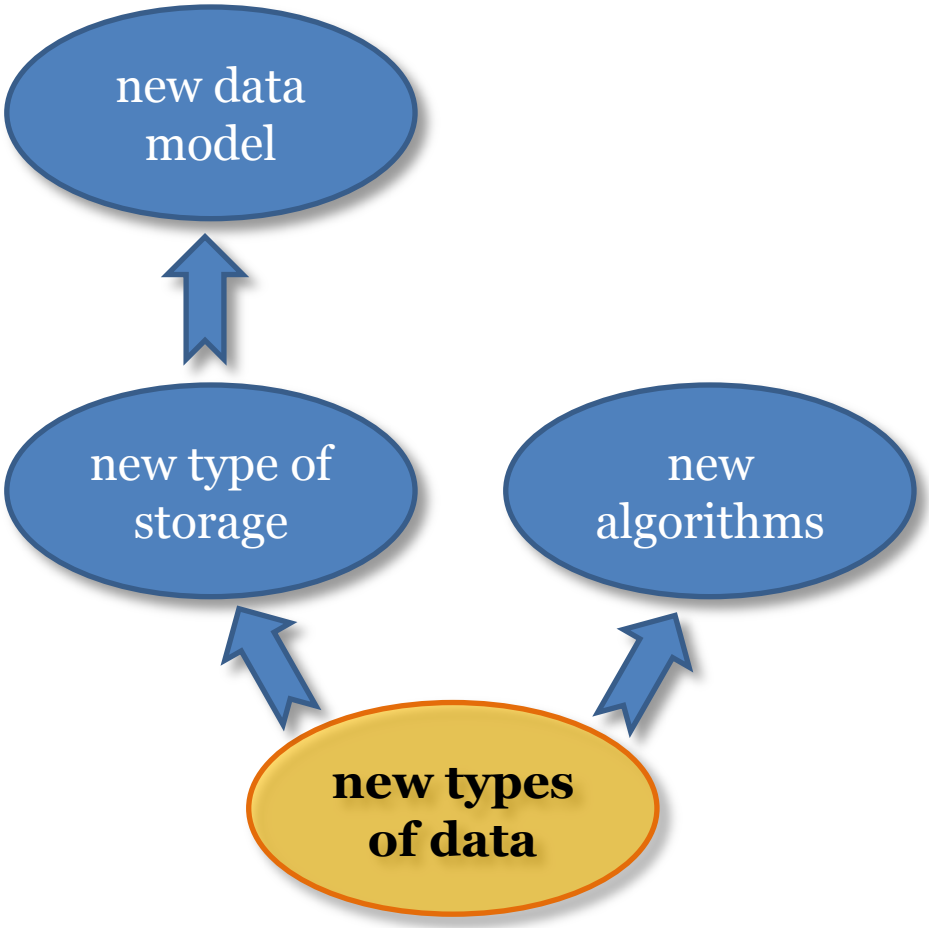

- geography data type
  - introduced in PostGIS 1.5 and MS SQL Server 2008
  - index over sphere
  - precise calculations over spheroid
  - slower calculations
  - still not in use as much as it should be
  - far fewer spatial functions

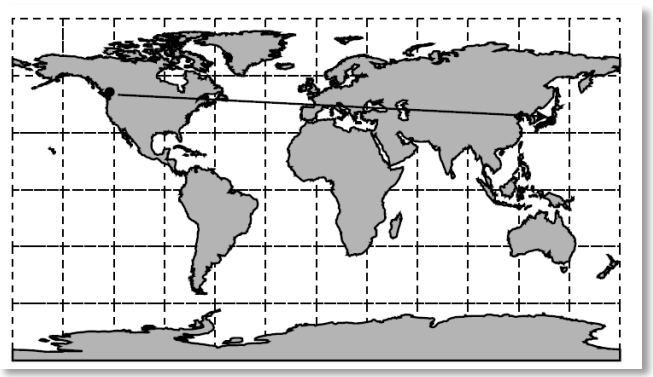

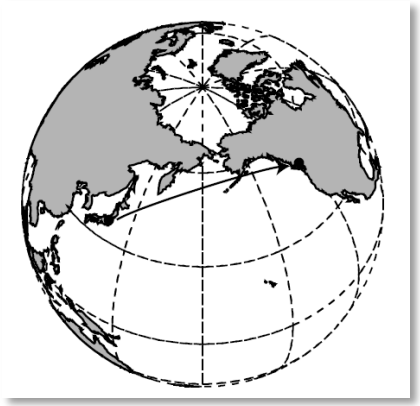

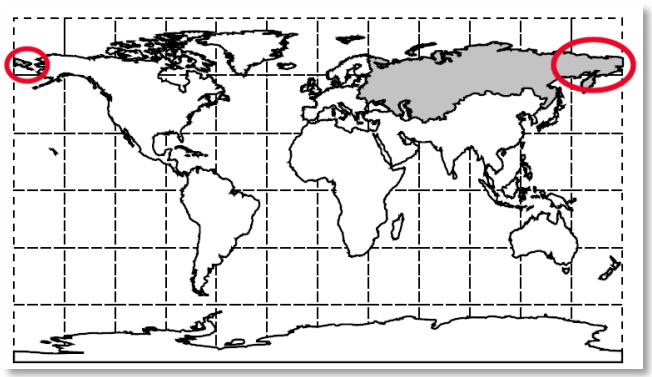

graphic: Alastair Aitchison,
Beginning Spatial with SQL Server 2008

- curves
  - introduced in PostGIS 1.2 and MS SQL Server “Denali”
  - CircularString (sequences of circular arcs)
  - CurvedPolygon (areas bounded by curved paths)
  - CompoundCurve (sequences of circular arcs and lines)
  - SQL/MM-3 standard (ISO/IEC 13249-3:2006)
  - fewer spatial functions then for LineStrings & Polygons

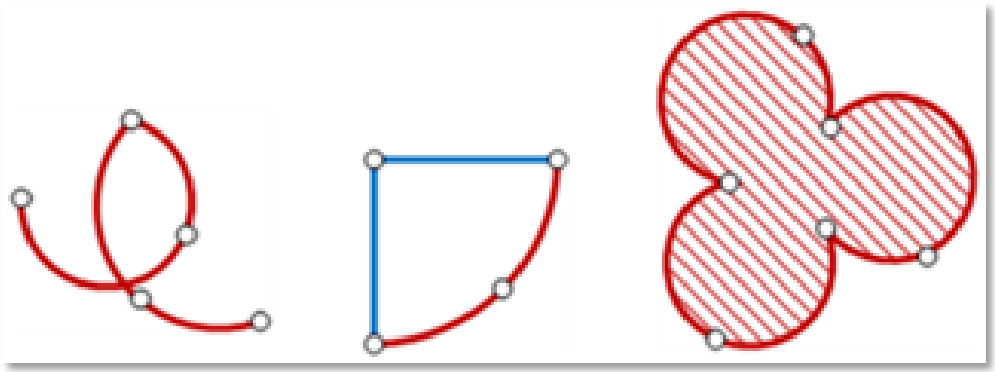

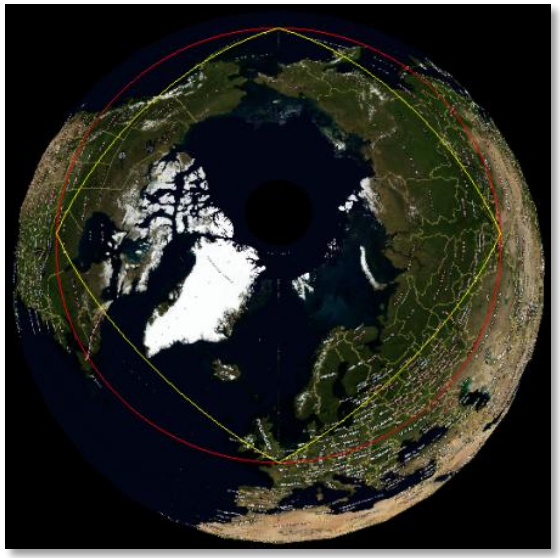

graphic: Ed Katibah , Milan Stojic ,
SQL Server Technical Article

- raster spatial data type
  - instead of raster data files
  - imported from jpegs, tiffs, pngs, digital elevation models
  - implemented in PostGIS 2.0 dev (as an add-on in 1.3.5+) and Oracle Spatial
  - big differences between two implementations
  - raster algebra
  - raster to vector and vector to raster conversion

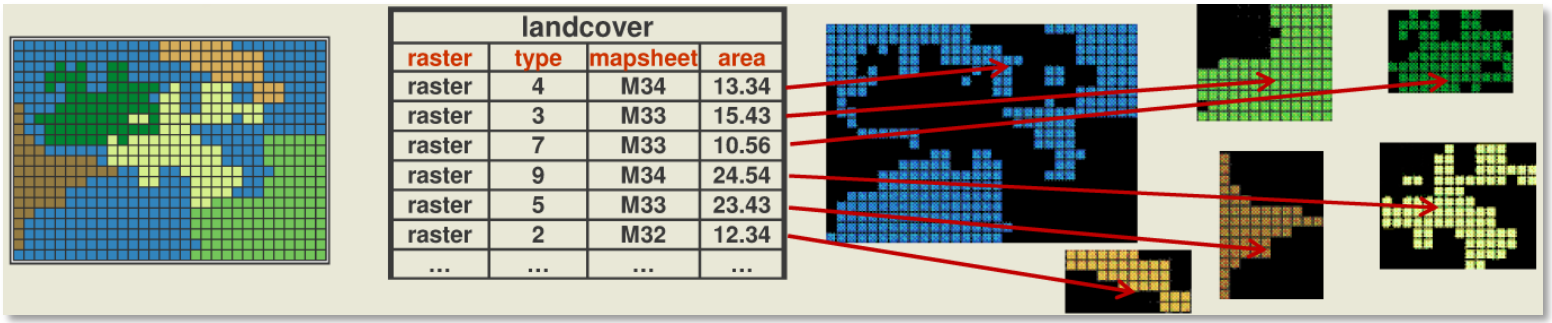

landcover

raster	type	mapsheet	area
raster	4	M34	13.34
raster	3	M33	15.43
raster	7	M33	10.56
raster	9	M34	24.54
raster	5	M33	23.43
raster	2	M32	12.34
...	...	...	...

graphic: Pierre Racine
Introduction to PostGIS WKT Raster and "Raster Objects"

Not Only SQL

"In computing, NoSQL is a term used to designate database management systems that differ from classic relational database management systems in some way. These data stores may not require fixed table schemas, and usually avoid join operations and typically scale horizontally."- *Wikipedia*

Key-value

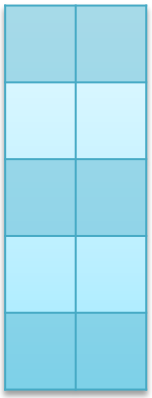

Bigtable

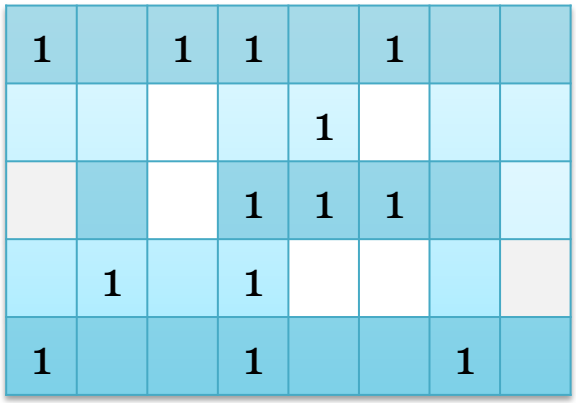

Document

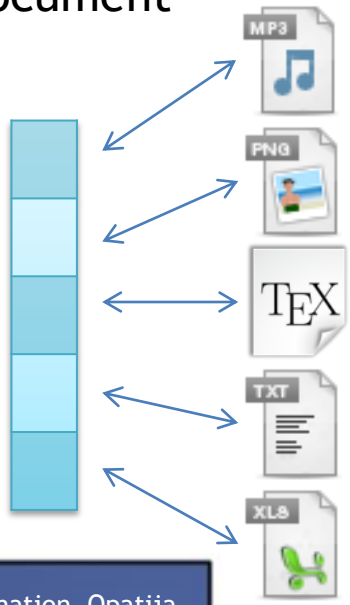

Graph DB

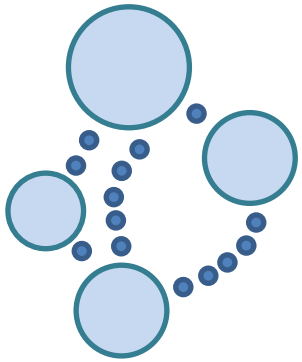

- NoSQL
- written in Erlang
- JSON storage
- exposes a RESTful HTTP API
- open source document-oriented database
- spatial extension for CouchDB
- R*Tree index
- Features:
  - (Multi-)Points
  - (Multi-)LineStrings
  - (Multi-)Polygons
- developed as a diploma thesis

typical 3-Tier architecture

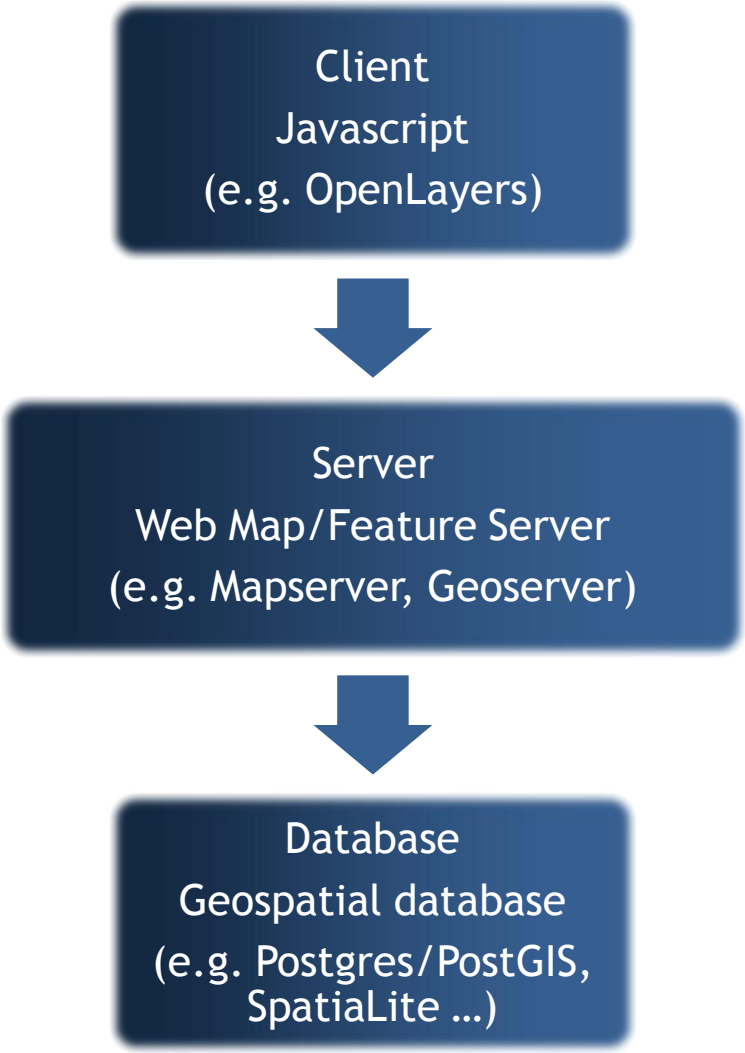

2-Tier architecture with CouchDB

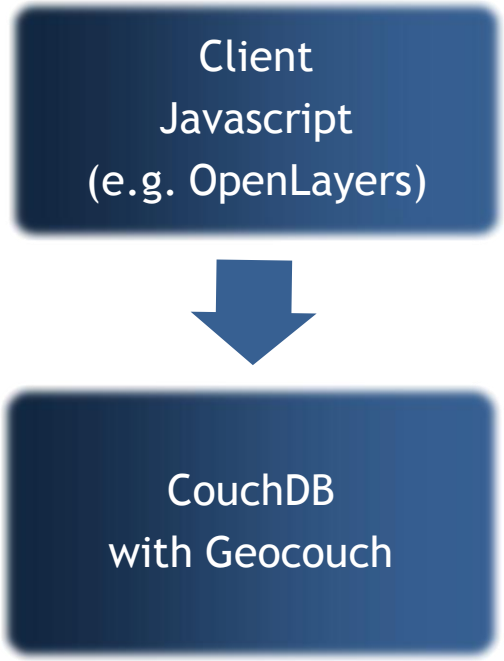

- NoSQL
- written in C
- BSON storage
- open source document-oriented database
- geohash spatial index (B-Tree)
  - points only
- geospatial query:
  - “find me the closest N items to my location”
  - “find me N items within …”

- NoSQL
- written in Java
- open source graph-oriented database
- model:
  - nodes
  - relationships between nodes
  - properties on both
- topology handling
- routing

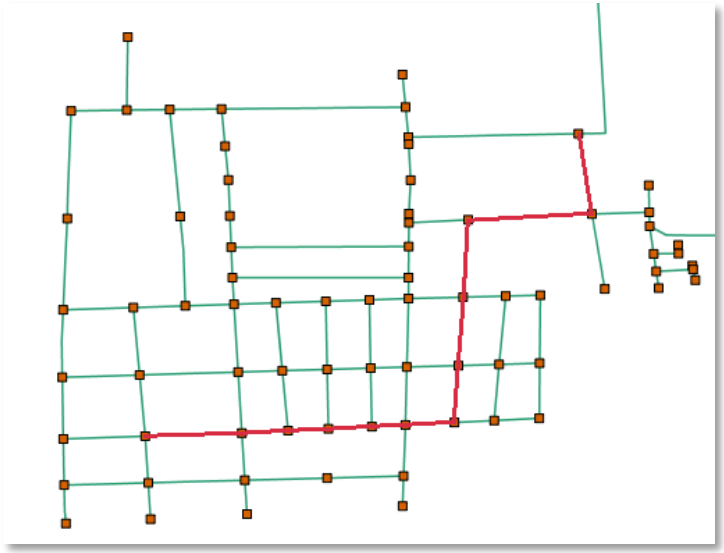

graphic: http://neotechnology.com/

Conclusion

- Storing raster as a raster data type in spatial databases will be a standard soon.
- NoSQL databases enable modeling and storing data specifically for a purpose.
- Graph NoSQL databases are the most promising for spatial data analysis.

Questions?