Spatial Databases: Lecture 10

Institute for Geoinformatics

Winter Semester 2014

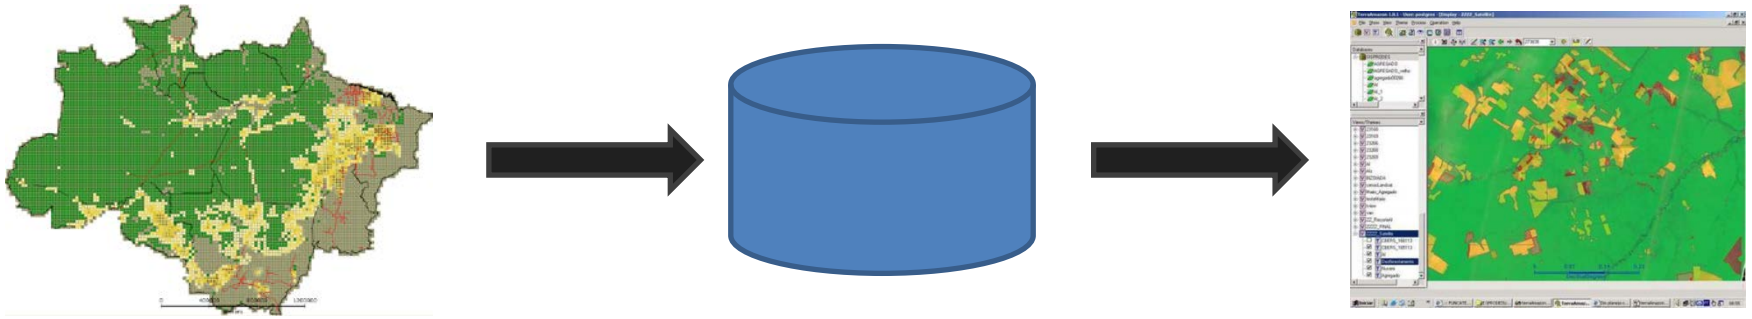

Malumbo Chipofya: room 109

SPATIAL INTELLIGENCE LABORATORY

Topic Overview

1. Prelude: Data and problem solving in science and applications
2. The Relational Database model
3. Interacting with relational databases
4. Spatial Relational Database Management Systems
5. Enlightenment: what is special about spatial - Prof. Dr. Gilberto Camara
6. **A sample of Nosql Databases: brief introductions + example applications**
   a. Array databases: SciDB
   b. Document databases: MongoDB
   c. Graph databases: Neo4J
7. Summary of all lectures given.

SPATIAL INTELLIGENCE LABORATORY

Alternative Database Technologies

- Relational Databases are here to stay!

Alternative Database Technologies

- Relational Databases are here to stay! But…
- Some criticism
  - Failure to scale at very high data volumes (longer tables -> longer query times)

Alternative Database Technologies

- Relational Databases are here to stay! But…
- Some criticism
  - Failure to scale at very high data volumes (longer tables -> longer query times)
  - Imposing a schema sometimes proves too rigid: many new applications are made possible by allowing more flexibility in data models

SPATIAL INTELLIGENCE LABORATORY

Alternative Database Technologies

- Relational Databases are here to stay! But…
- Some criticism
  - Failure to scale at very high data volumes (longer tables -> longer query times)
  - Imposing a schema sometimes proves too rigid: many new applications are made possible by allowing more flexibility in data models
  - Complex relationships difficult to represent explicitly – conflicts between the need for normalization and the complexity of the data model result in **join bombs**

SPATIAL INTELLIGENCE LABORATORY

Alternative Database Technologies

- Relational Databases are here to stay! But…
- Some criticism: The **Join Bomb**

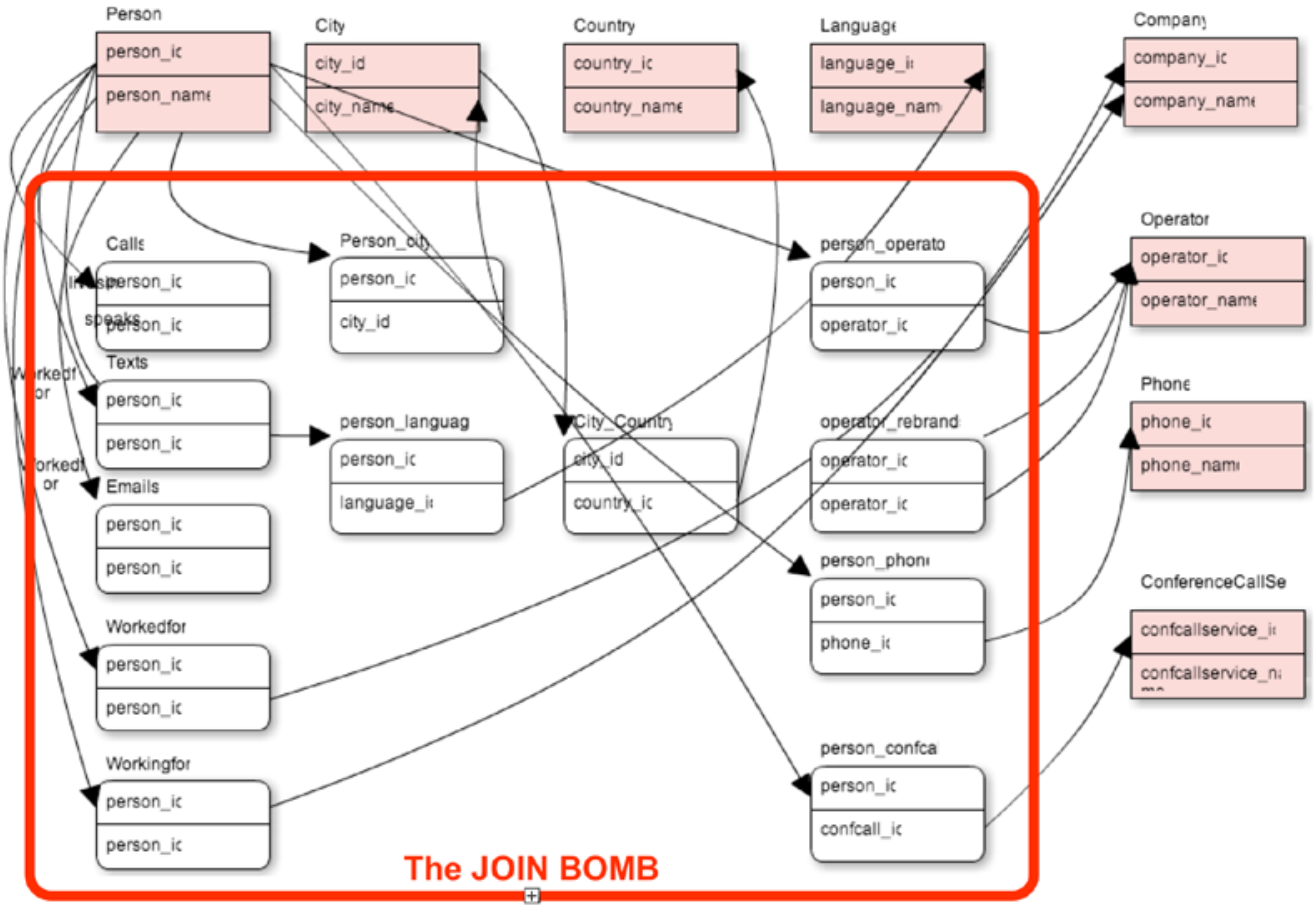

Image: **Rik Van Bruggen.** Learning Neo4j. Packt Publishing, Birmingham, UK, 2014. pg 33.

SPATIAL INTELLIGENCE LABORATORY

Alternative Database Technologies

- Relational Databases are here to stay! But…
- Some alternatives

Key-value store

Column Family stores

Document stores

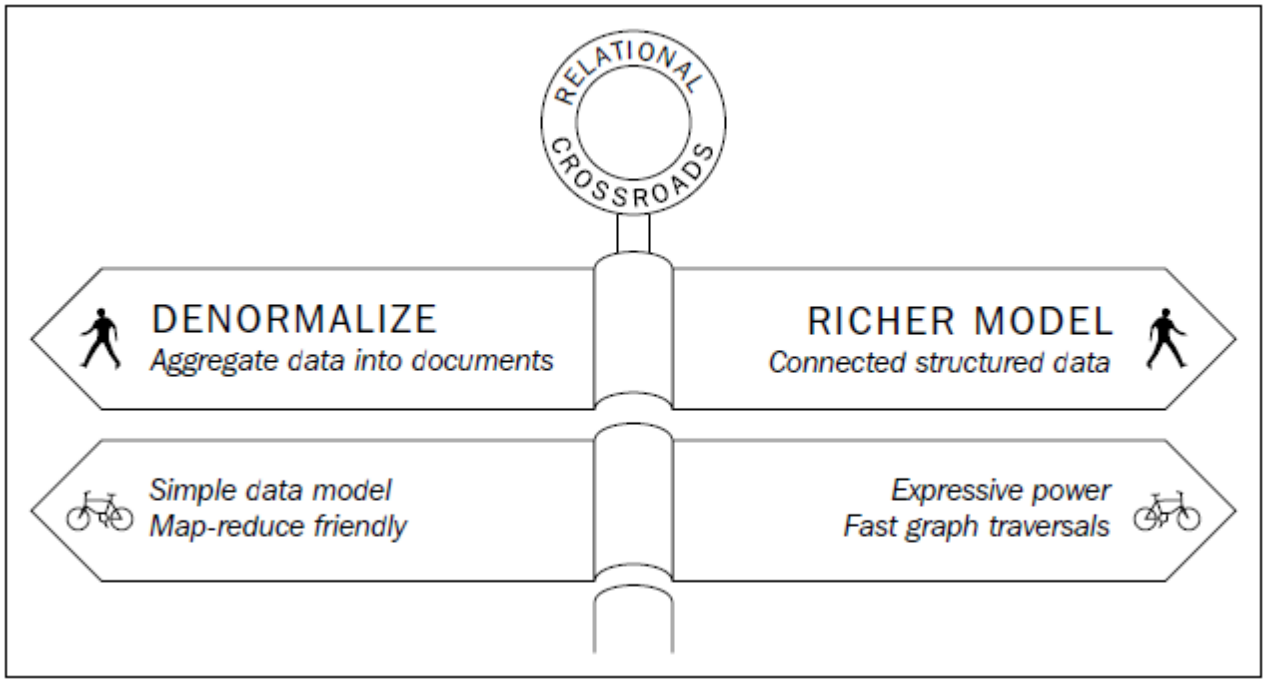

Graph Databases

Van Bruggen. pg 33.

SPATIAL INTELLIGENCE LABORATORY

Graphs

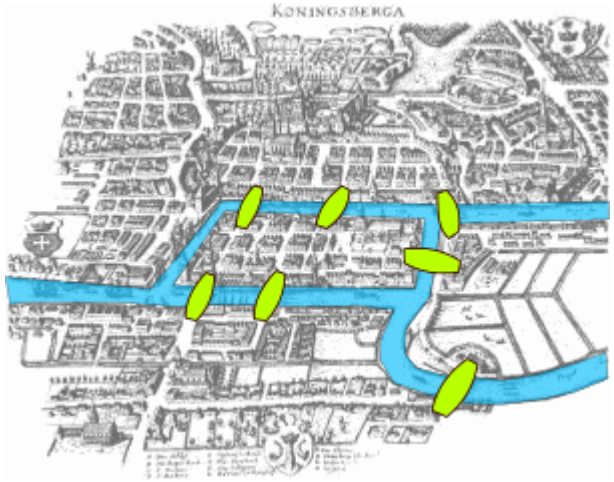

SPATIAL INTELLIGENCE LABORATORY

Graphs

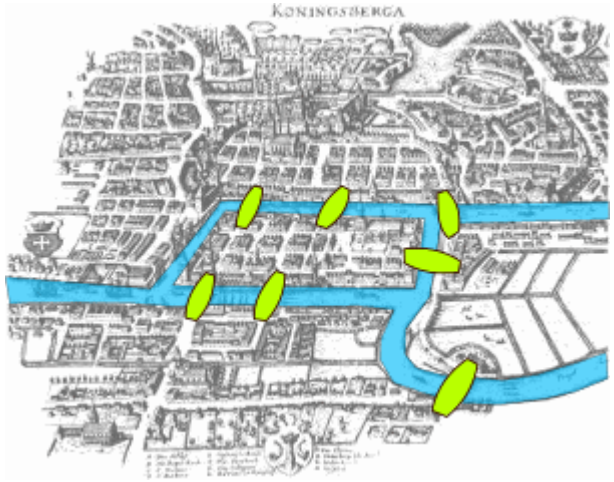

The 7 Bridges of Königsberg problem

- Cross every bridge exactly once

SPATIAL INTELLIGENCE LABORATORY

Graphs

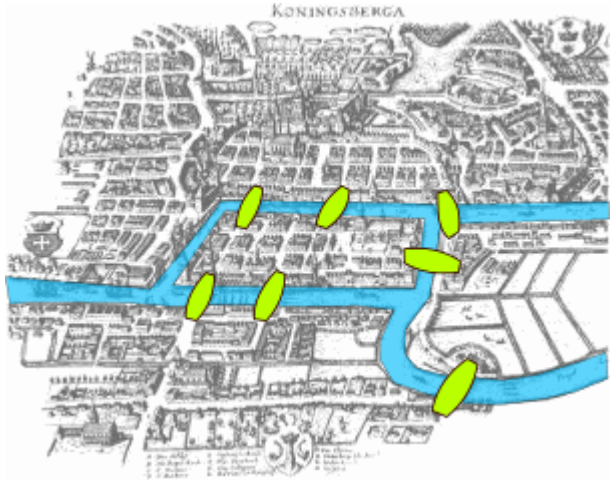

The 7 Bridges of Königsberg problem

- Cross every bridge exactly once
- Solved by Leonhard Euler in 1735 (pub. 1736)

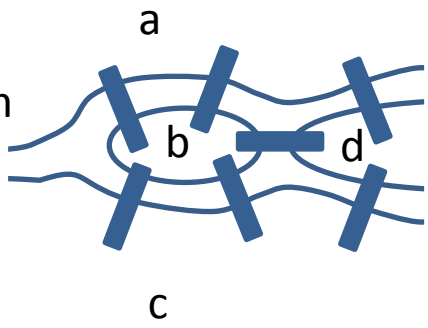

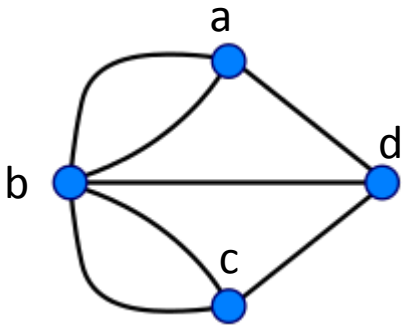

Images (top & bottom):
http://en.wikipedia.org/wiki/Seven_Bridges_of_Königsberg

SPATIAL INTELLIGENCE LABORATORY

Graphs

- Nodes

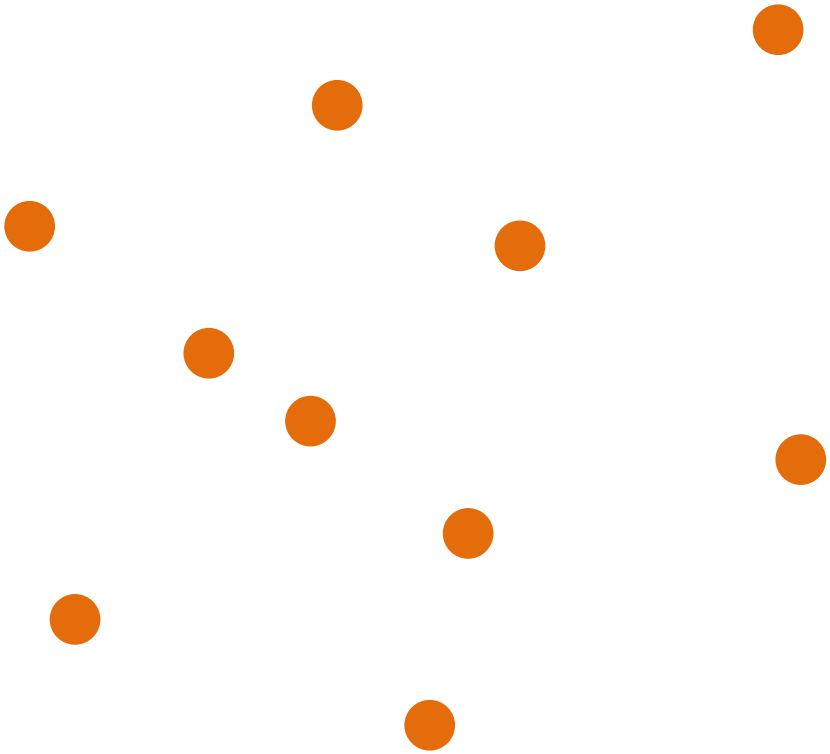

SPATIAL INTELLIGENCE LABORATORY

Graphs

- Nodes
- Edges

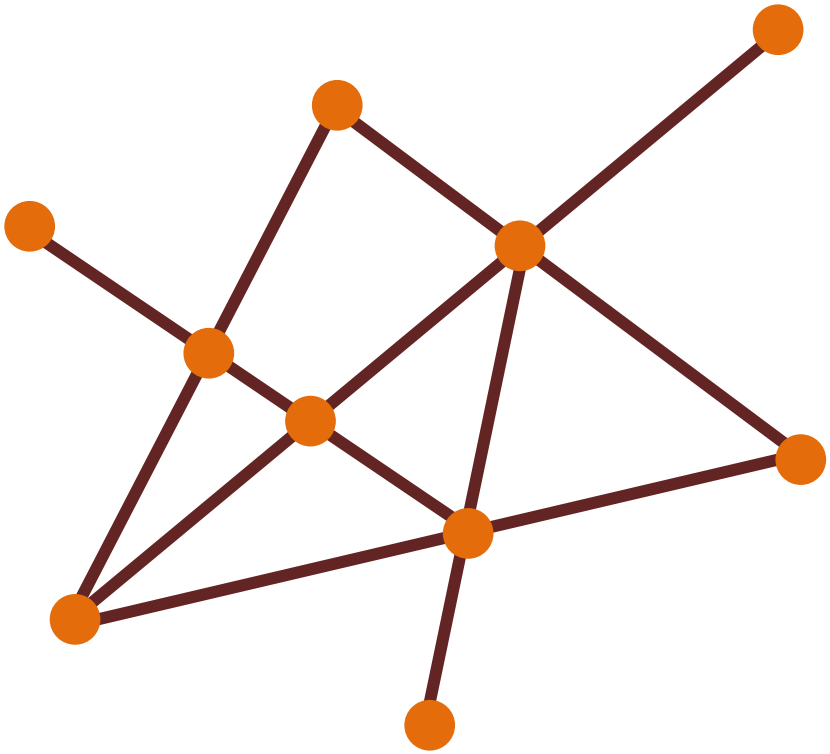

SPATIAL INTELLIGENCE LABORATORY

Graphs

- Nodes
- Edges
- Attributes

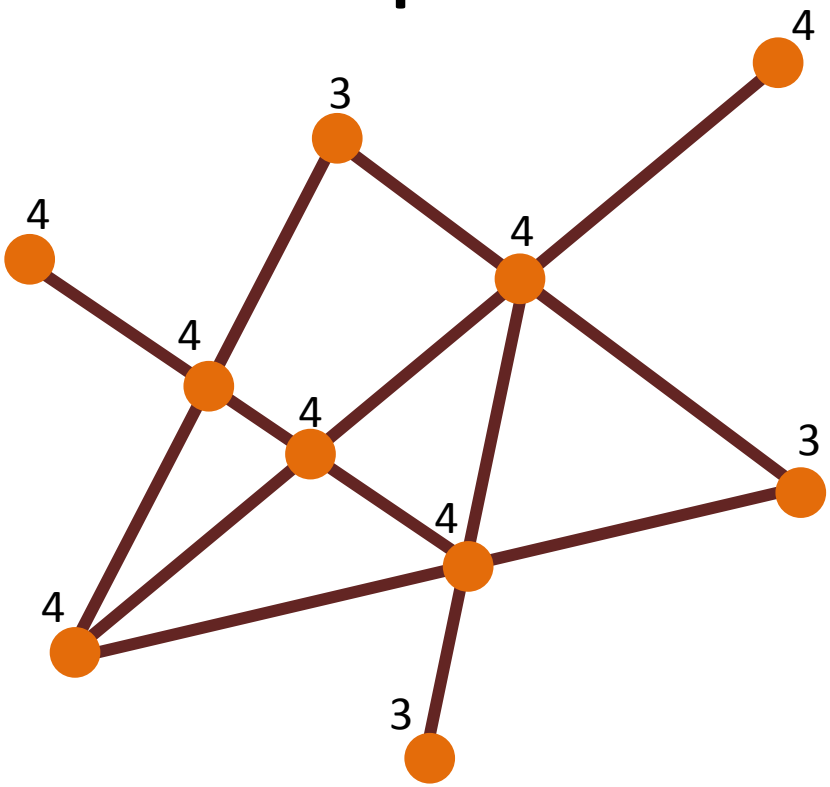

SPATIAL INTELLIGENCE LABORATORY

Graphs

- Nodes
- Edges

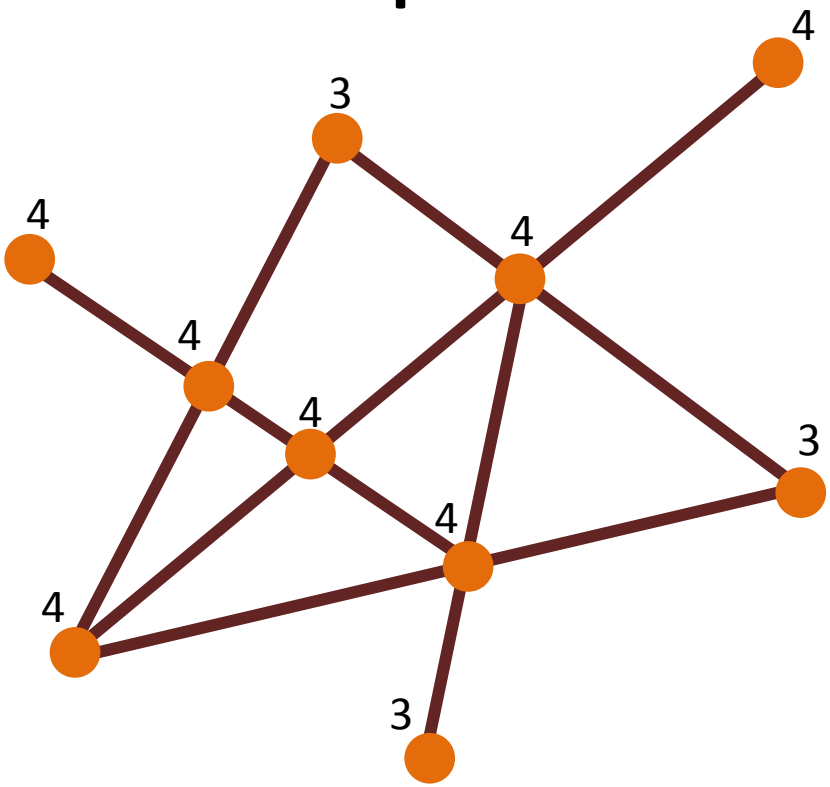

- Attributes: number of points in longest collinear set containing the point

SPATIAL INTELLIGENCE LABORATORY

Applications of Graphs

- Used as representation for a vast majority of computational tasks.
- They are general and yet can be made very precise.

Applications of Graphs

- Modelling relationships between places and resources (Natural gas flows network – http://enipedia.tudelft.nl/)

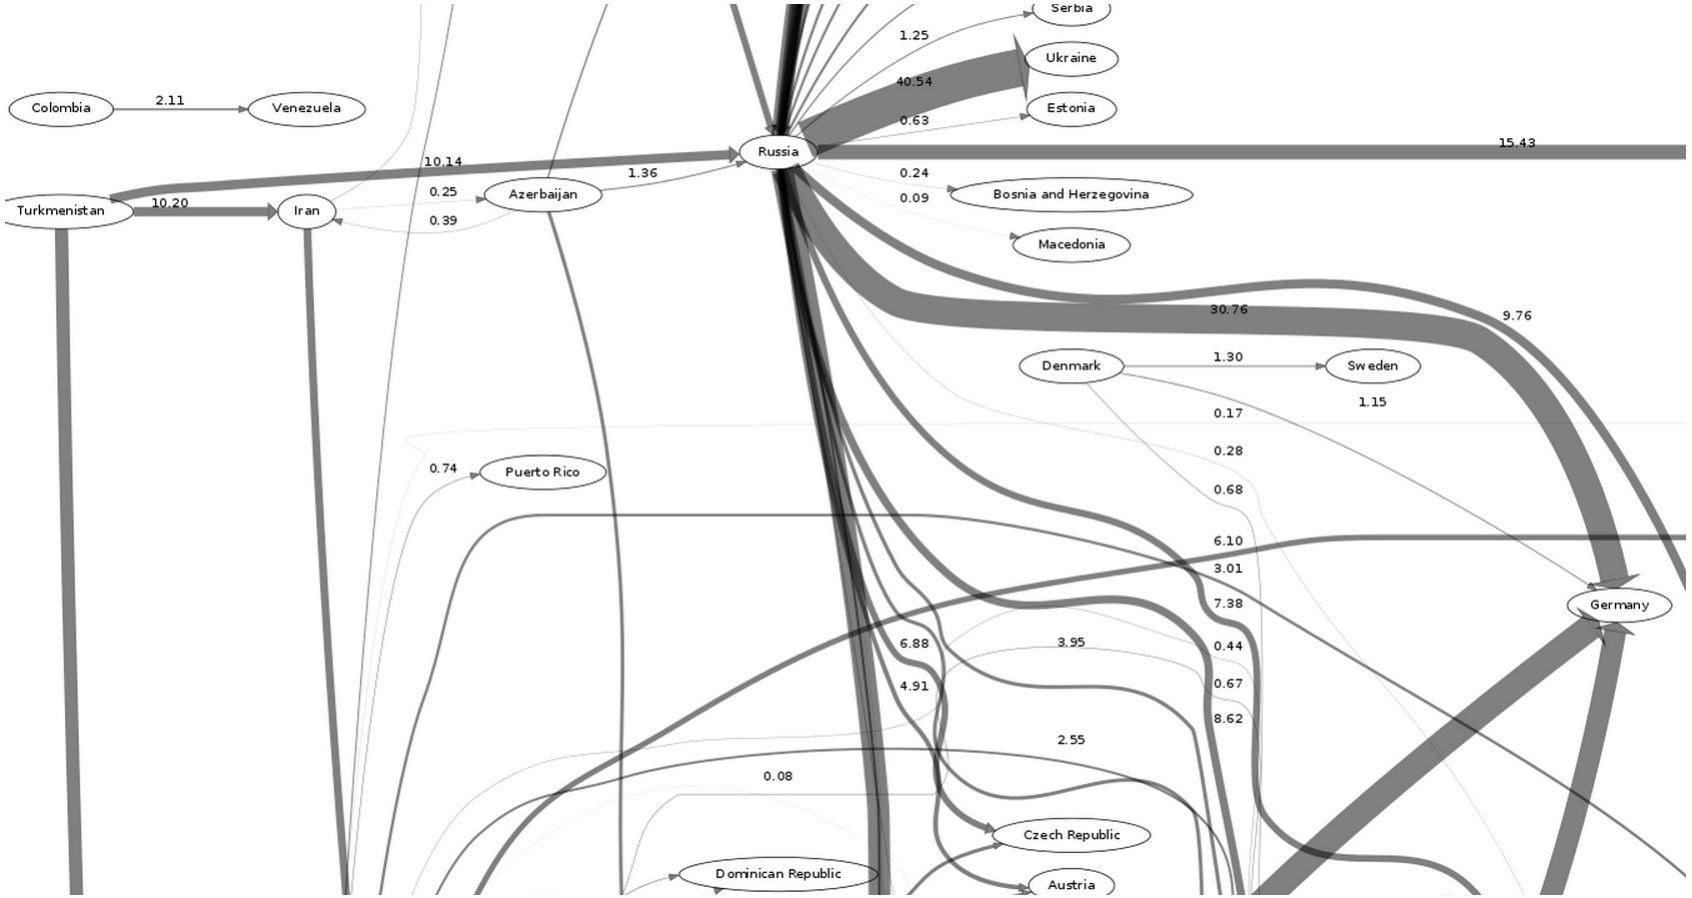

Applications of Graphs

- Route finding (remember pgRouting's Dijkstra and A*?

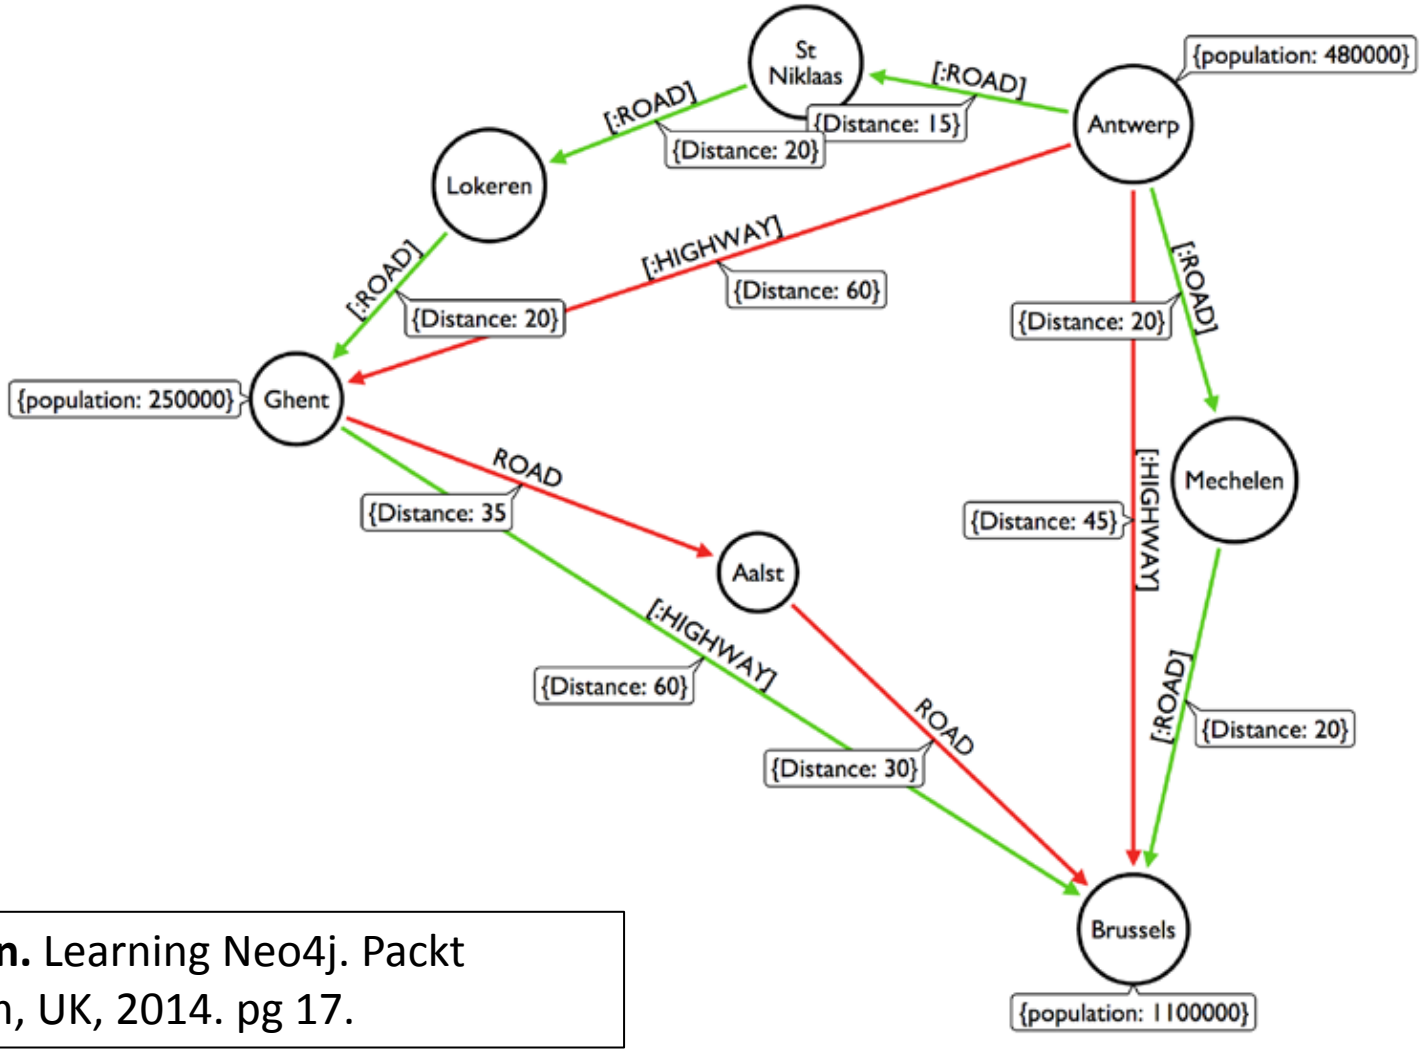

Image: **Rik Van Bruggen.** Learning Neo4j. Packt Publishing, Birmingham, UK, 2014. pg 17.

Applications of Graphs

- Recommendations in social and business applications: you may also like/know/need/etc.

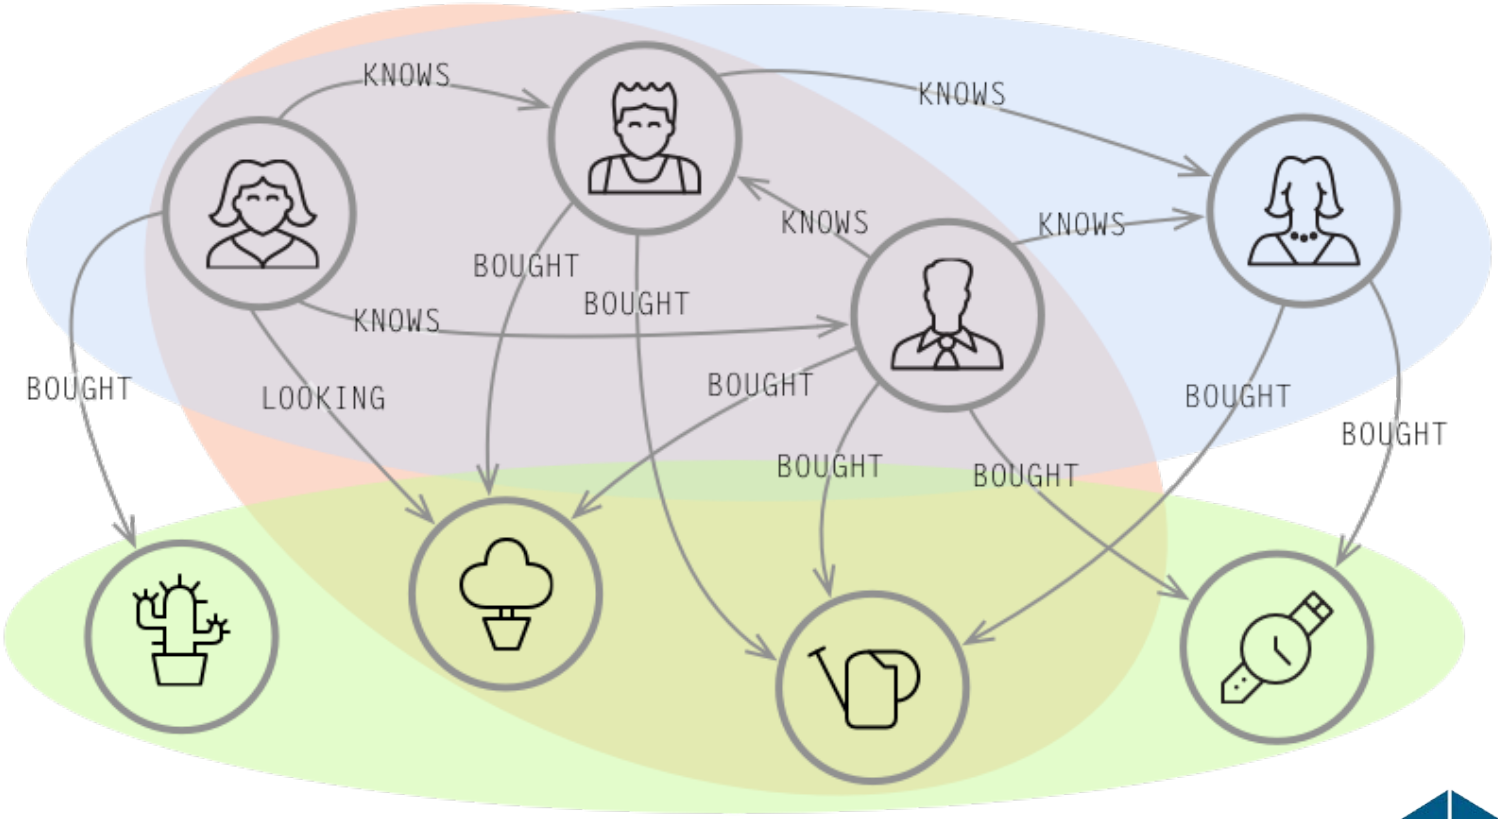

Image: http://neo4j.com/use-cases

SPATIAL INTELLIGENCE LABORATORY

The Property Graph Data Model

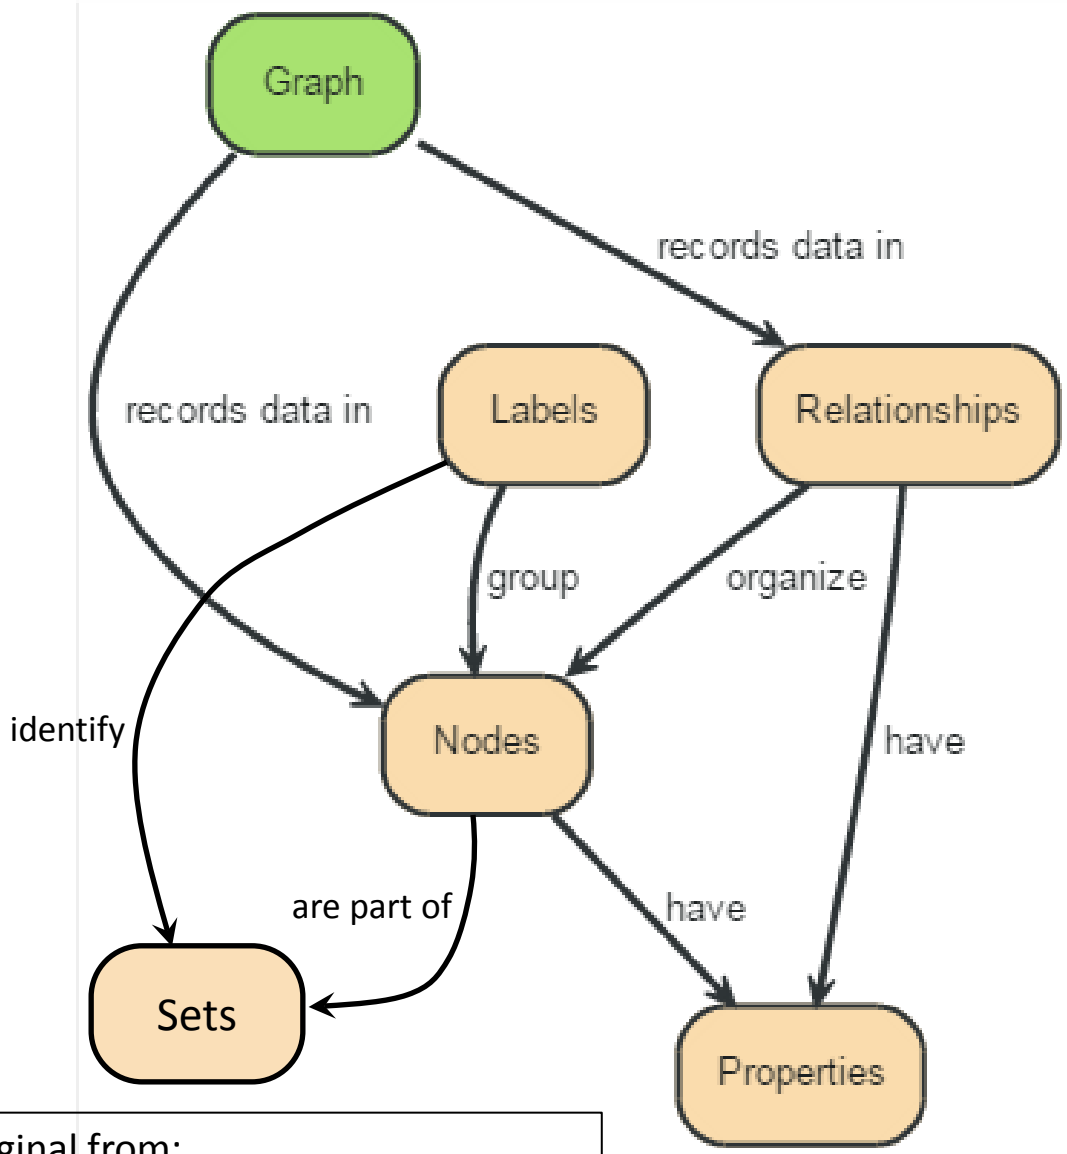

Image modified. Original from:
http://neo4j.com/docs/2.1.6/what-is-a-graphdb.html

SPATIAL INTELLIGENCE LABORATORY

The Property Graph Data Model

Example instance? Anyone?

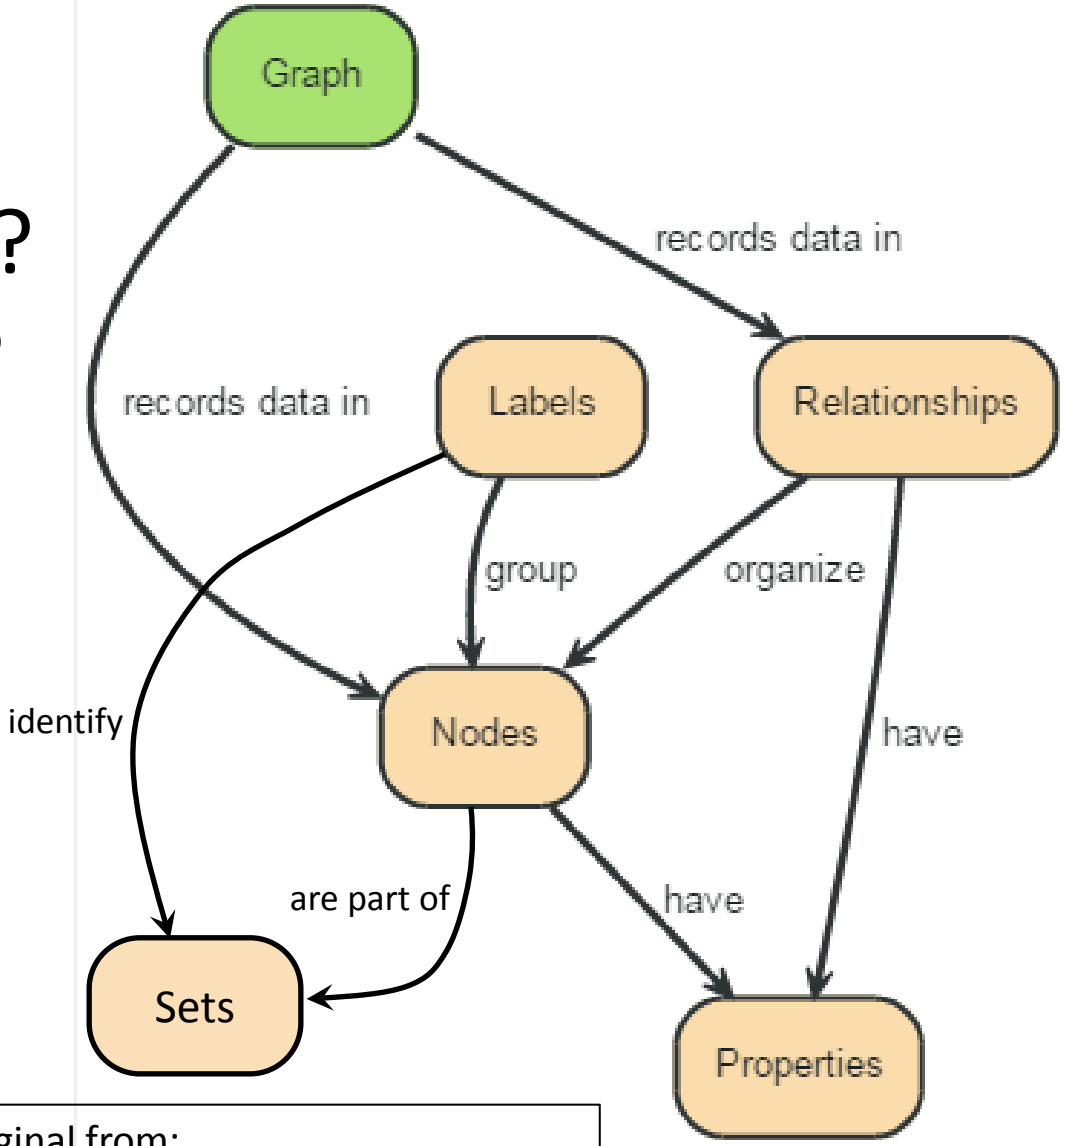

Image modified. Original from:
http://neo4j.com/docs/2.1.6/what-is-a-graphdb.html

SPATIAL INTELLIGENCE LABORATORY

The Property Graph Data Model

- We've already seen another example

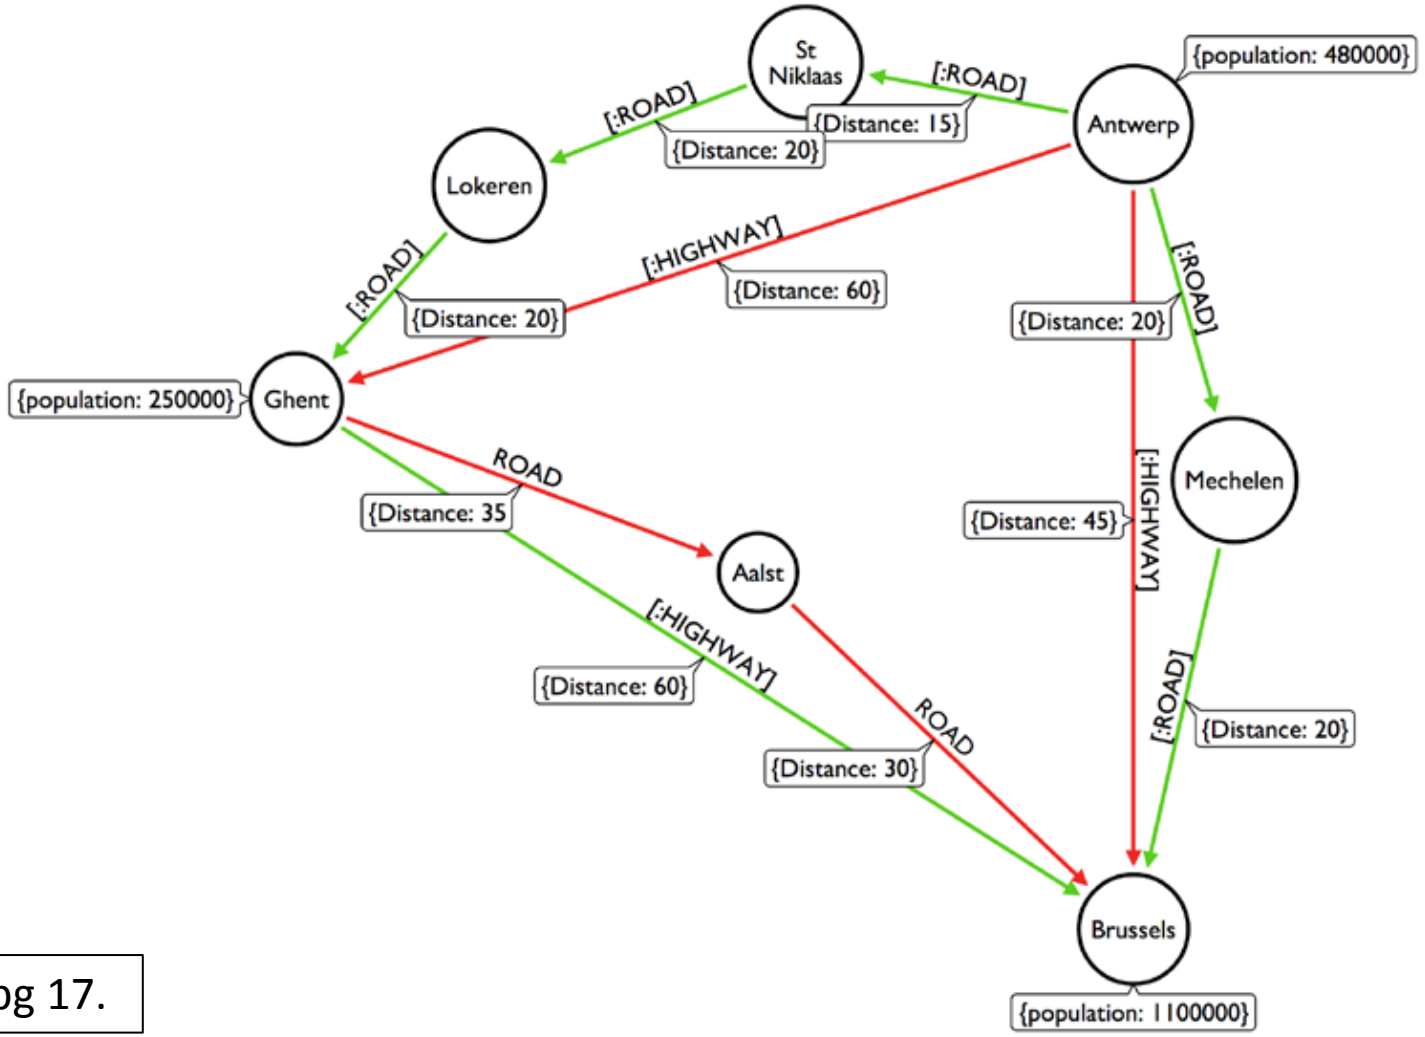

Image: **Van Bruggen.** pg 17.

The Property Graph Data Model

- It is suited for directed labelled multirelational graphs
- Another example

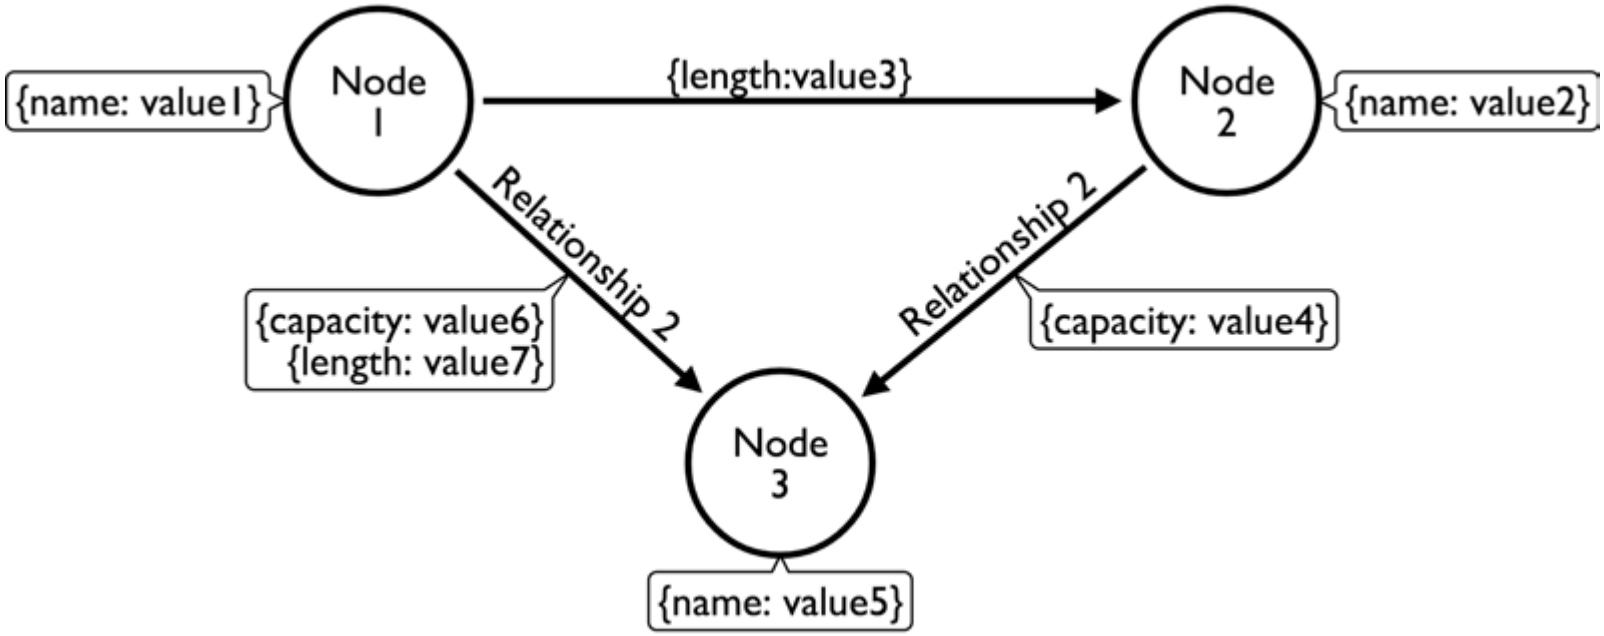

Image: **Van Bruggen.** pg 35.

(Property) Graph Database System

- Stores (property) graphs natively: on disk, the data are represented directly as graphs
- And provides all necessary data management
  - Indexing
  - Constraint specification
  - Querying
  - Transaction management
  - Security
  - Etc.

SPATIAL INTELLIGENCE LABORATORY

(Property) Graph Database System

- Querying done by graph traversal

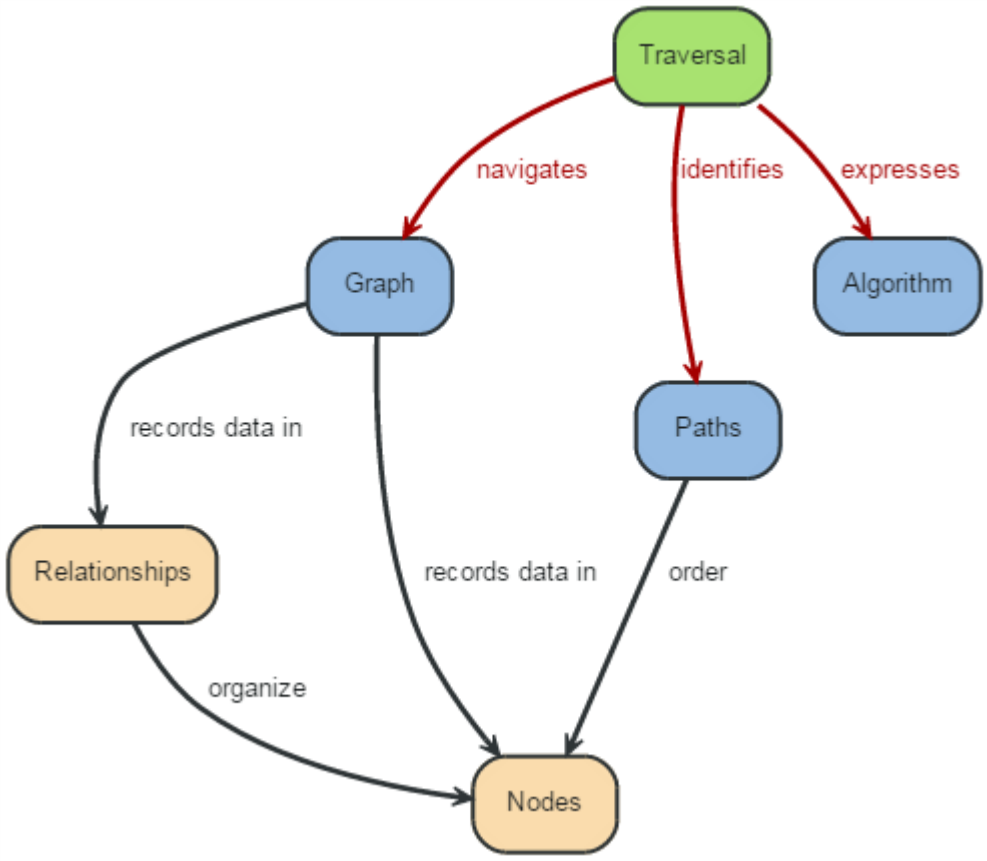

SPATIAL INTELLIGENCE LABORATORY

(Property) Graph Database System

- Indexing node and relationship properties speeds up query processing

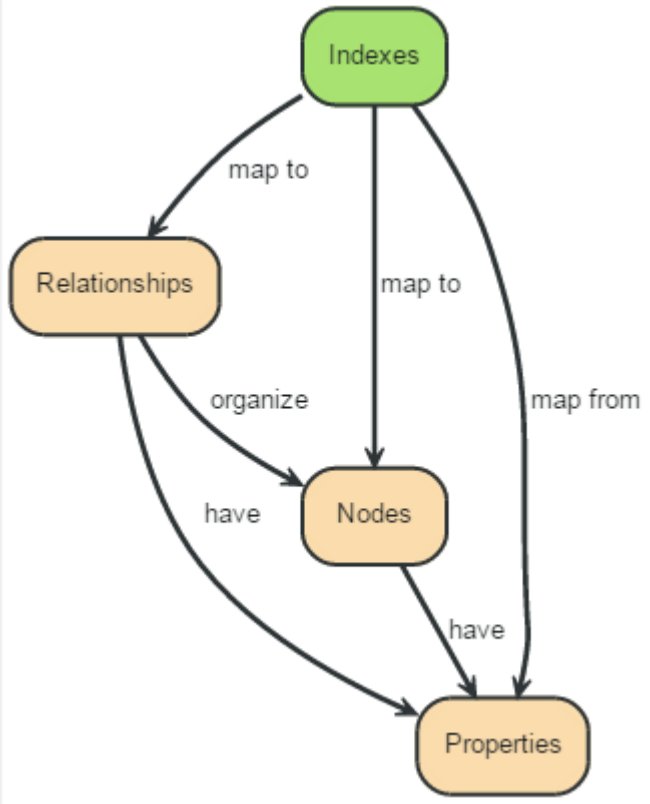

SPATIAL INTELLIGENCE LABORATORY

(Property) Graph Database System

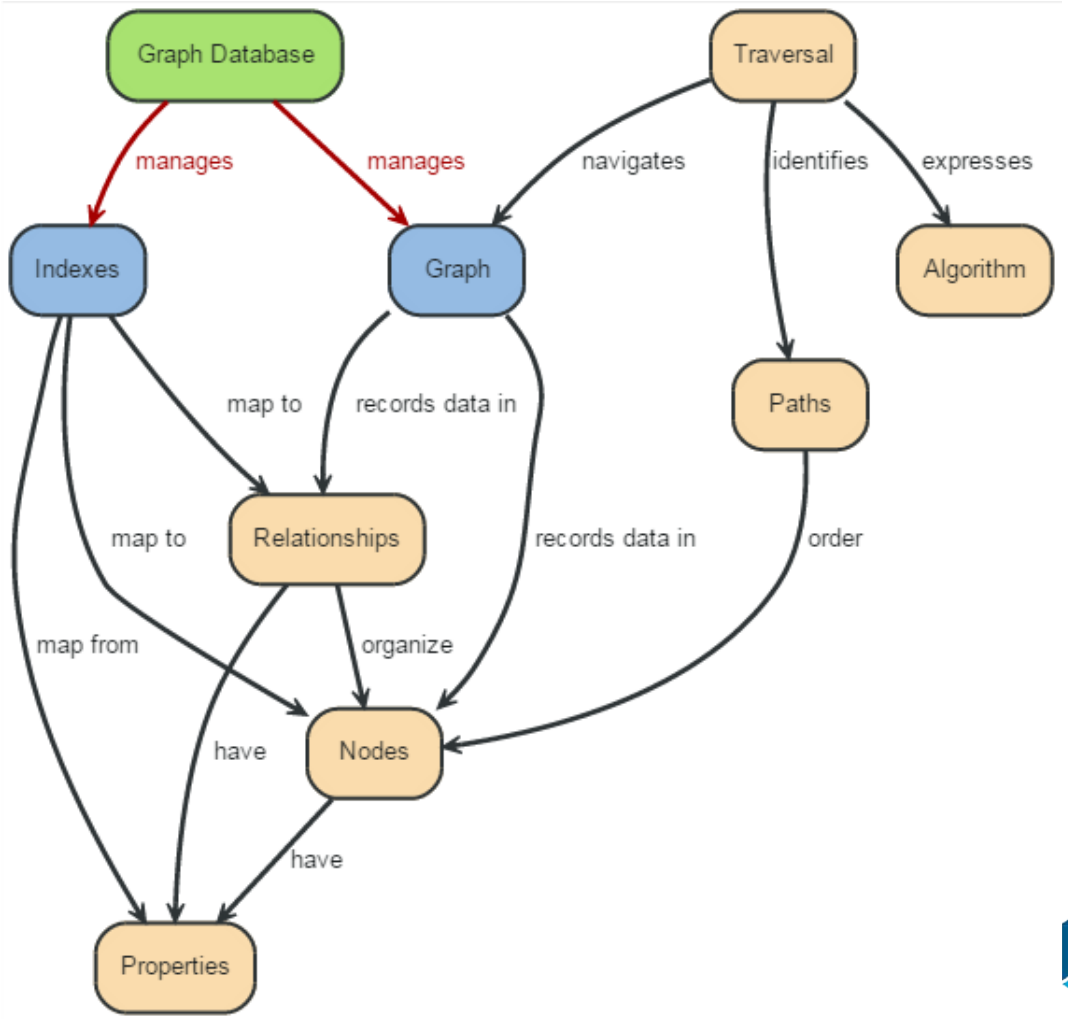

SPATIAL INTELLIGENCE LABORATORY

(Property) Graph Database System

- Neo4j is an example of a graph database
  - Stores property graphs (of course, pfuff)
  - Provides a declarative query language: Cypher
  - Can be accessed in various ways including via a RESTFUL API (http)
  - Comes with a built-in web-based GUI – the graph browser which supports visualization
  - Is open source – hence lots of plugins (including spatial ones are beginning to appear)

Neo4j: An introduction

- Data Model Elements

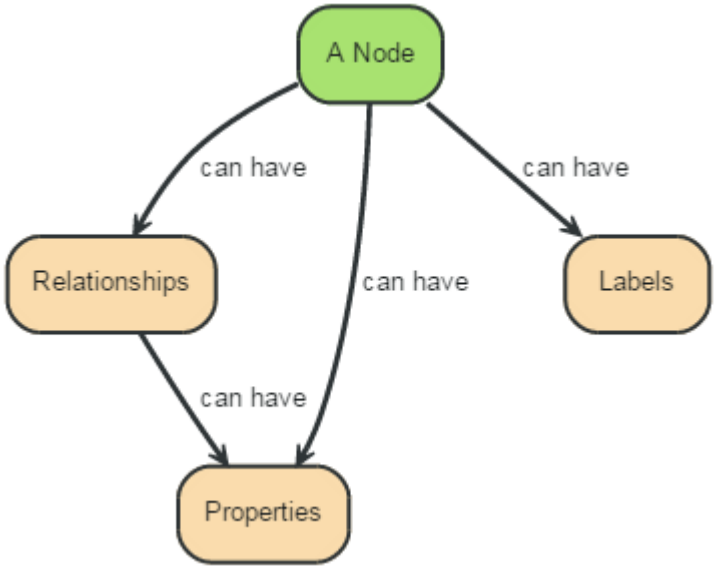

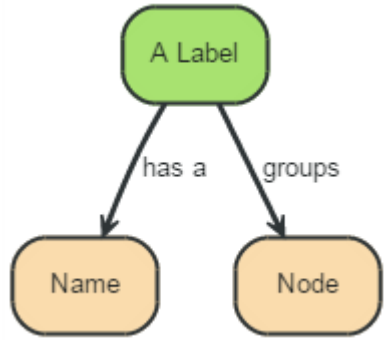

- Examples:
  - Property: name
  - Values: Maja, Alice, Oscar, William

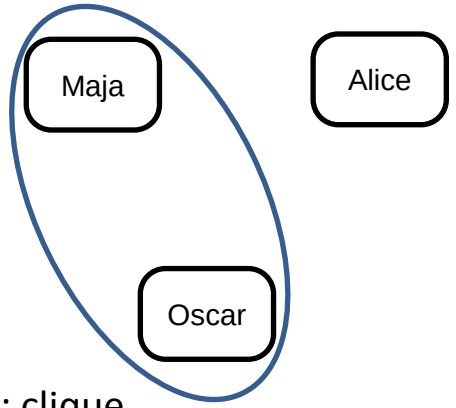

  - Label: clique
  - Targets: Maja, Oscar

SPATIAL INTELLIGENCE LABORATORY

Neo4j: An introduction

- Data Model Elements

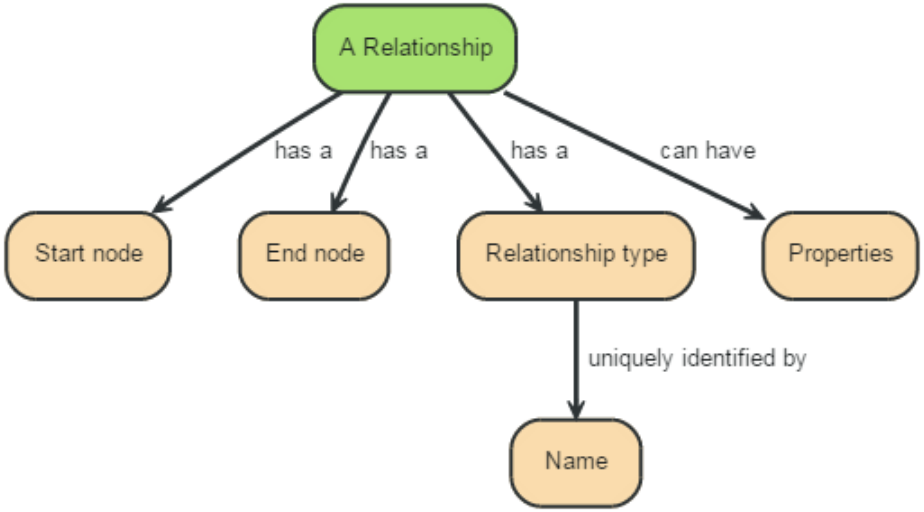

- Examples

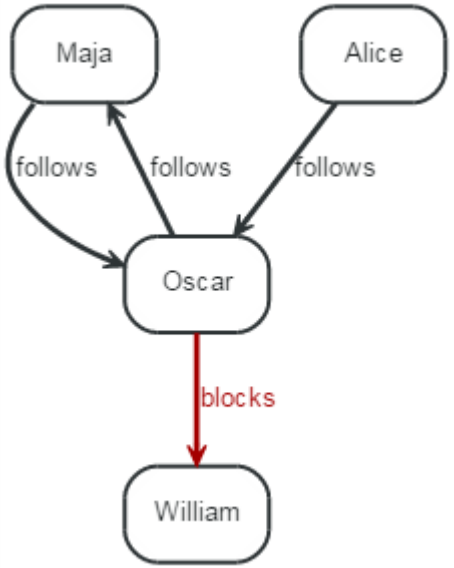

SPATIAL INTELLIGENCE LABORATORY

Neo4j: An introduction

- Data Model Elements

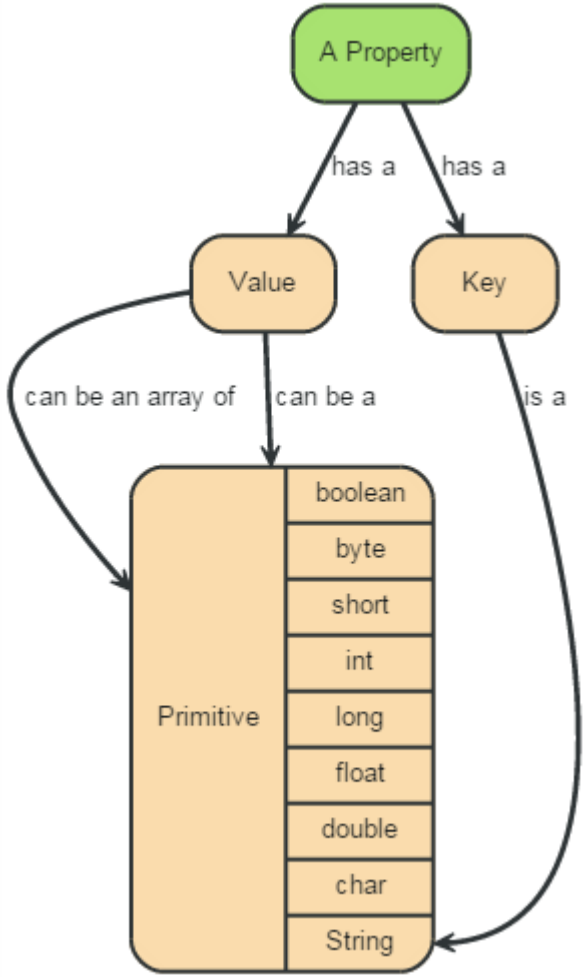

- Examples

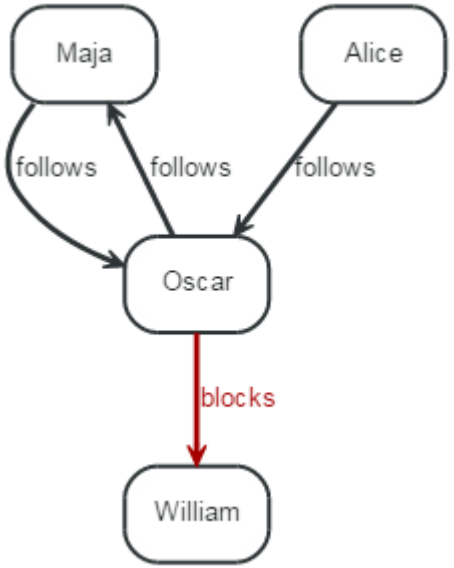

SPATIAL INTELLIGENCE LABORATORY

Neo4j: An introduction

- Data Model Elements

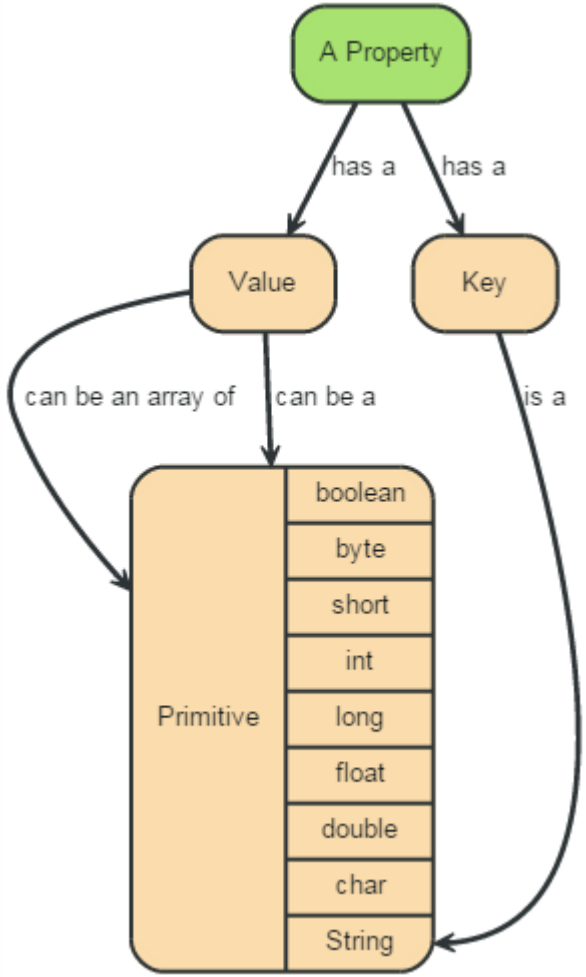

- Examples

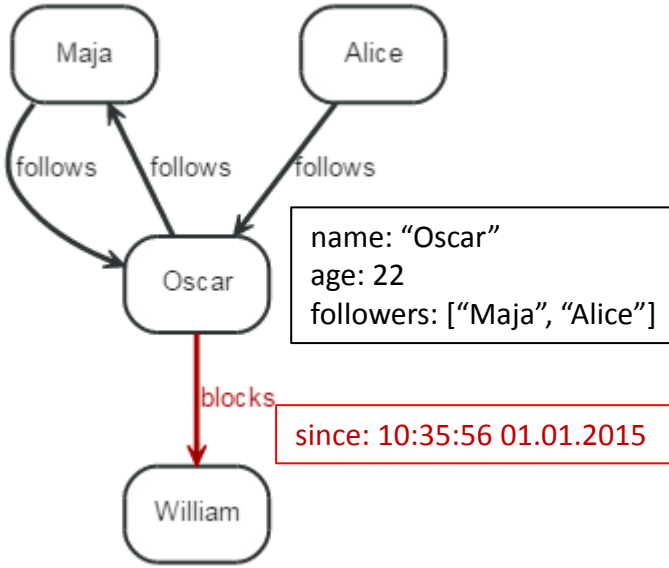

SPATIAL INTELLIGENCE LABORATORY

Neo4j: An introduction

- Data Model Elements

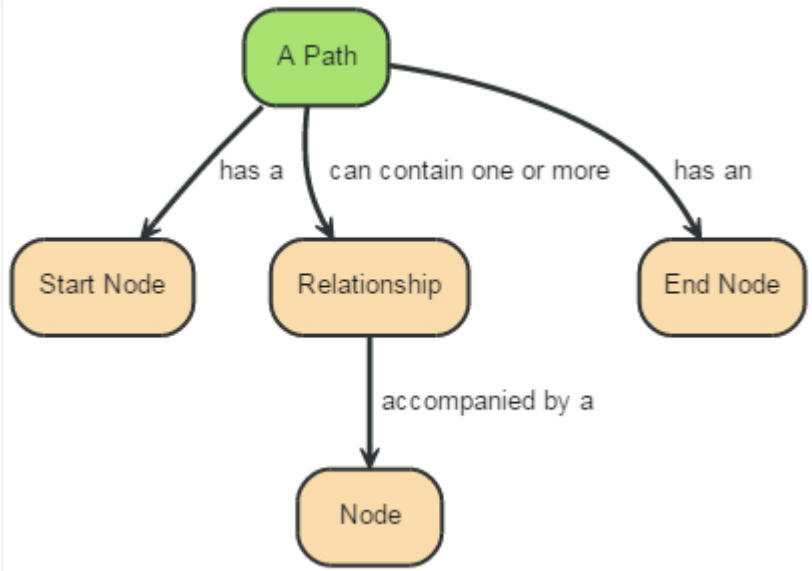

- Examples
  - Alice -> Oscar -> William

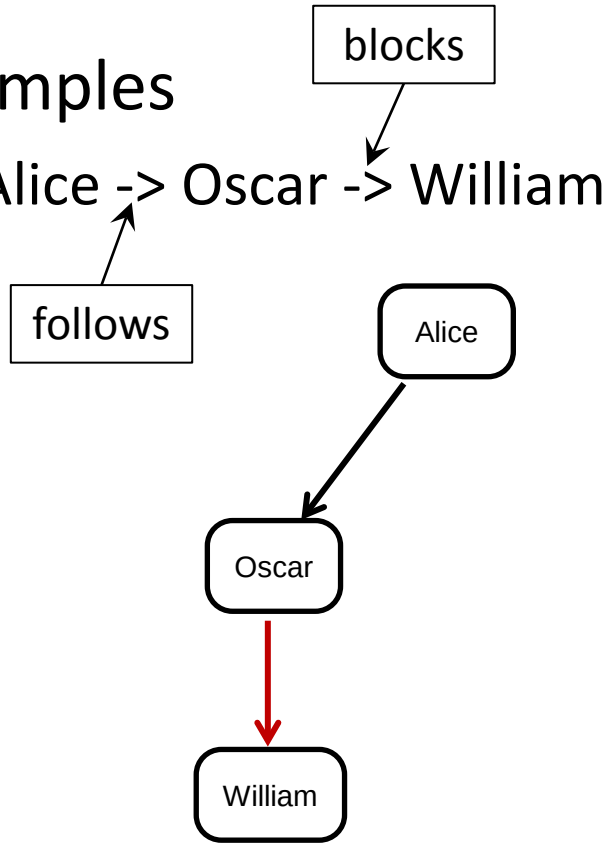

SPATIAL INTELLIGENCE LABORATORY

Neo4j: An introduction

- Cypher PATTERNS
  - A node: (n)

Neo4j: An introduction

- Cypher PATTERNS
  - A node: (n)
  - Related nodes: (n)-->(m)<--()--(a)

SPATIAL INTELLIGENCE LABORATORY

Neo4j: An introduction

- Cypher PATTERNS
  - A node: (n)
  - Related nodes: (n)-->(m)<--()--(a)
  - Labels: (n:Number)-->(m:Moles)

SPATIAL INTELLIGENCE LABORATORY

Neo4j: An introduction

- Cypher PATTERNS
  - A node: (n)
  - Related nodes: (n)-->(m)<--()--(a)
  - Labels: (n:Number)-->(m:Moles)
  - Naming relationships: (a)-[r]->(b)

SPATIAL INTELLIGENCE LABORATORY

Neo4j: An introduction

Identifiers: n, m, a, b, r

- Cypher PATTERNS
  - A node: (n)
  - Related nodes: (n)-->(m)<--()--(a)
  - Labels: (n:Number)-->(m:Moles)
  - Naming relationships: (a)-[r]->(b)
  - Typing relationships: (a)-[r:Follows]->(b)

SPATIAL INTELLIGENCE LABORATORY

Neo4j: An introduction

Identifiers: n, m, a, b, r

- Cypher PATTERNS
  - A node: (n)
  - Related nodes: (n)-->(m)<--()--(a)
  - Labels: (n:Number)-->(m:Moles)
  - Naming relationships: (a)-[r]->(b)
  - Typing relationships: (a)-[r:Follows]->(b)
  - Properties
    - On nodes: (p {name: “Malu", hobby: “Eating" })
    - On relationships: (a)-[{blocked: false}]->(b)

SPATIAL INTELLIGENCE LABORATORY

Neo4j: An introduction

Identifiers: n, m, a, b, r

- Cypher PATTERNS
  - A node: (n)
  - Related nodes: (n)-->(m)<--()--(a)
  - Labels: (n:Number)-->(m:Moles)
  - Naming relationships: (a)-[r]->(b)
  - Typing relationships: (a)-[r:Has]->(b)-[:Was]->(n)
  - Properties
    - On nodes: (p {name: “Malu", hobby: “Eating" })
    - On relationships: (a)-[{blocked: false}]->(b)
  - Paths: (a)-[*2]->(b)-[*2..3]->(n)-[*]->(b)-[*..3]->

SPATIAL INTELLIGENCE LABORATORY

Neo4j: An introduction

Identifiers: n, m, a, b, r

- Cypher PATTERNS
  - A node: (n)
  - Related nodes: (n)-->(m)<--()--(a)
  - Labels: (n:Number)-->(m:Moles)
  - Naming relationships: (a)-[r]->(b)
  - Typing relationships: (a)-[r:Has]->(b)-[:Was]->(n)
  - Properties
    - On nodes: (p {name: “Malu", hobby: “Eating" })
    - On relationships: (a)-[{blocked: false}]->(b)
  - Paths: (a)-[*2]->(b)-[*2..3]->(n)-[*]->(b)-[*..3]->

SPATIAL INTELLIGENCE LABORATORY

Neo4j: An introduction

- Go to

Start menu > Programs > Neo4j Community > Neo4j Community

- You'll see something like this

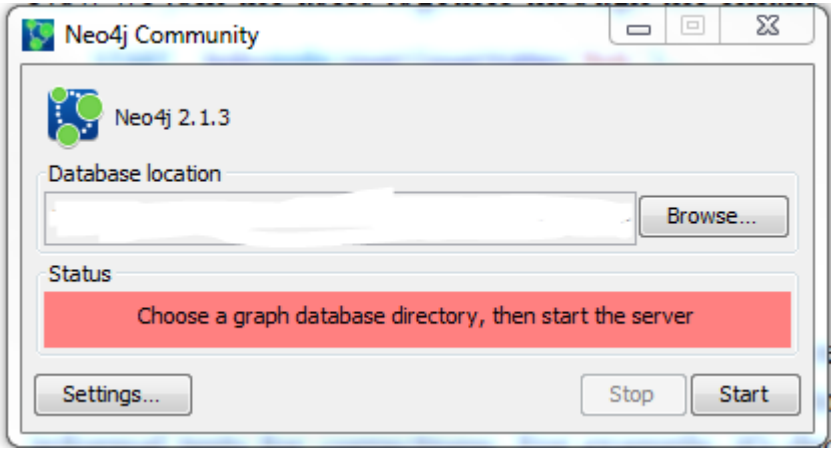

- Click start

SPATIAL INTELLIGENCE LABORATORY

Neo4j: An introduction

- Go to

Start menu > Programs > Neo4j Community > Neo4j Community

- You'll see something like this

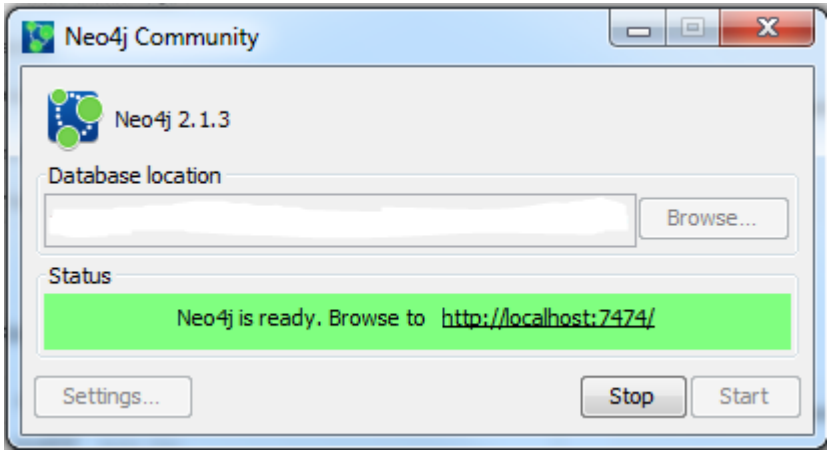

- Follow the link and play around

SPATIAL INTELLIGENCE LABORATORY

Neo4j: An introduction

- CREATE
- MATCH
- RETURN
- WITH
- LIMIT, SKIP
- MERGE
- Etc…

SPATIAL INTELLIGENCE LABORATORY

Neo4j: An introduction

- A Cypher query has a structure similar to an SQL one – let's see how to
  - Create nodes and relationships with CREATE
  - Query with MATCH
  - Update the graph
  - Traverse (find a path in) the graph

References

- **Rik Van Bruggen.** Learning Neo4j. Packt Publishing, Birmingham, UK, 2014.
- **Ian Robinson, Jim Webber, and Emil Eifrem.** Graph Databases. O'Reilly Media, Sebastopol, USA, 2013.
- http://neo4j.com/docs/2.1.6/ - 31/12/2014

SPATIAL INTELLIGENCE LABORATORY

That’s all for today

Thank you!

Questions?

SPATIAL INTELLIGENCE LABORATORY