



Save your relation with a Graph



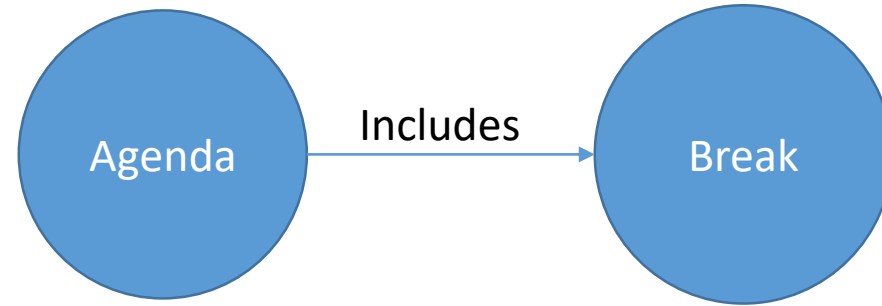
Graph databases in .NET

Agenda

- Introduction
- Use cases
- Tools and first code

break

- Modelling
- Graph languages
- CRUD
- Querying

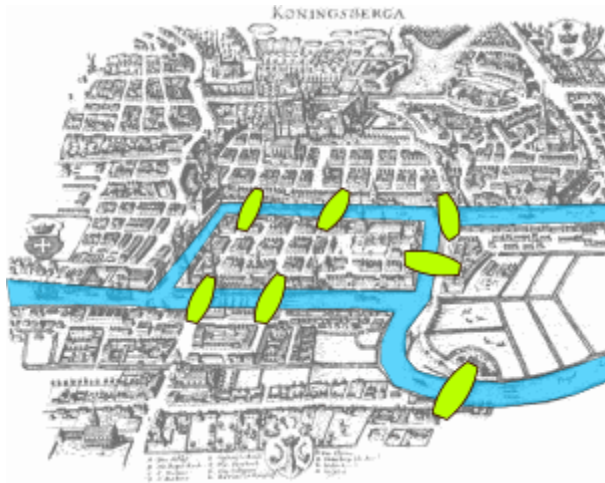


Introduction

Graph theory

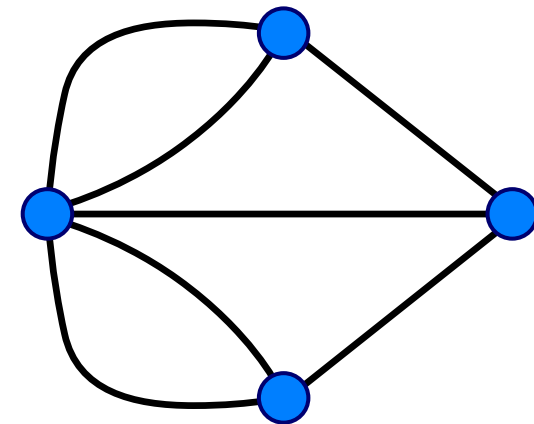
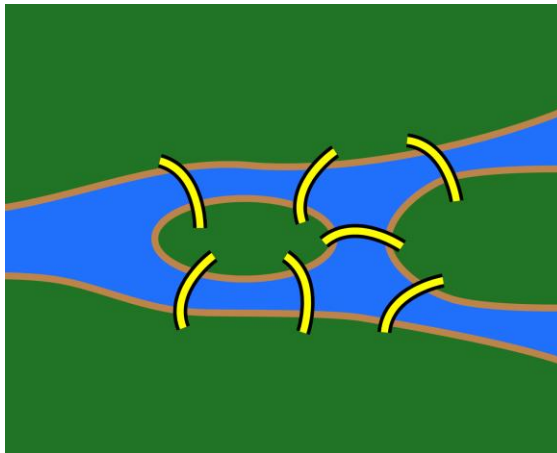
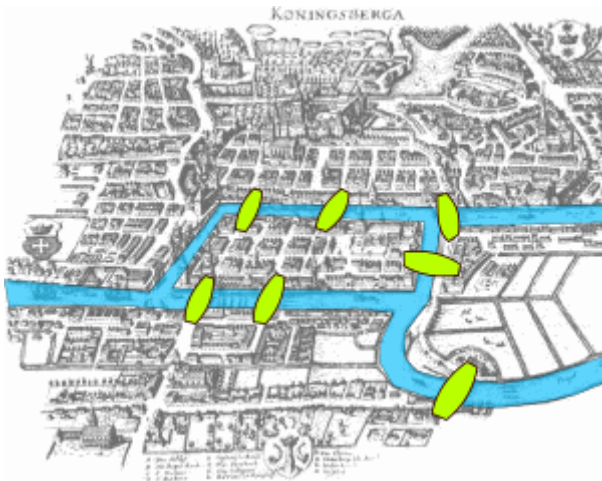
First paper by Leonhard Euler (1707-1783)

Seven Bridges of Königsberg, published in 1736



Seven Bridges of Königsberg

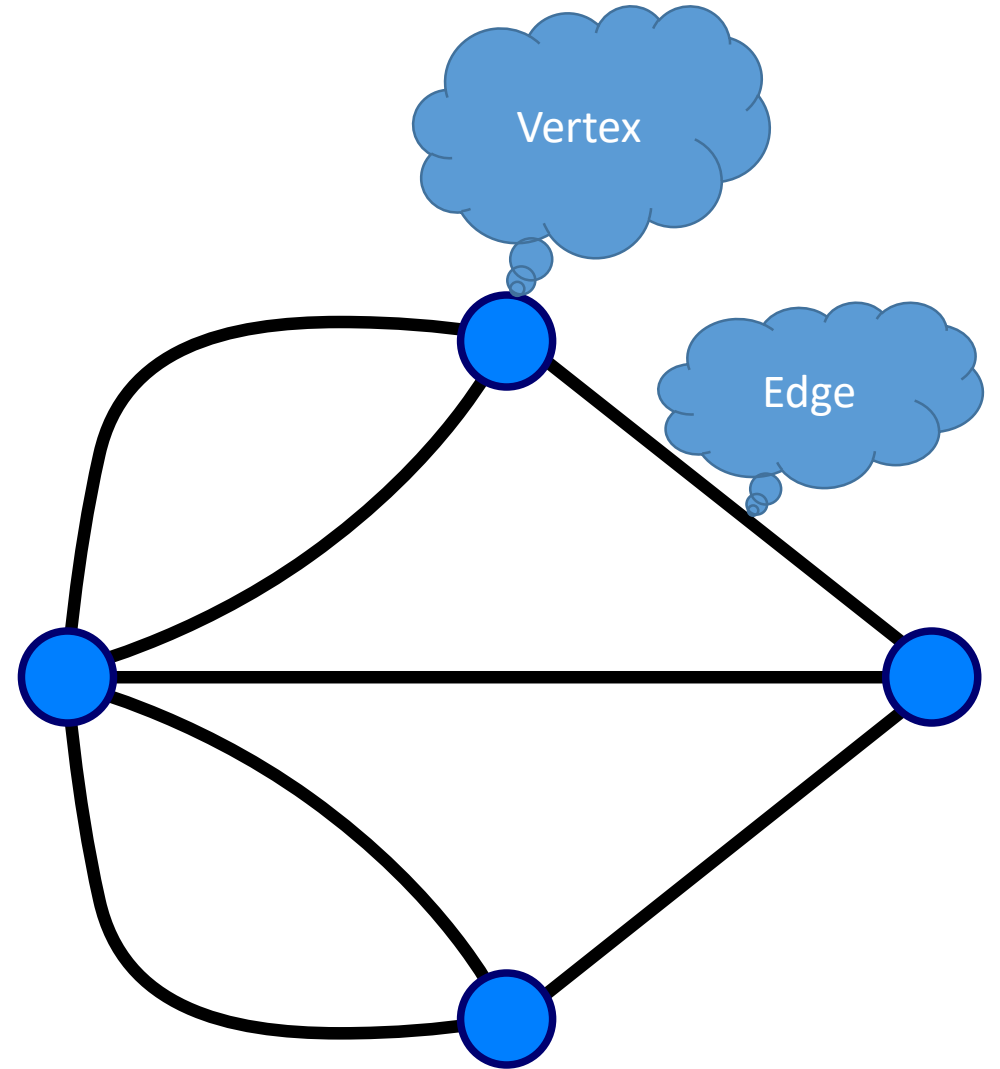
Devise a walk through the city that would cross each of those bridges once and only once.



Vertices and Edges

Or nodes and relationships...

Or nodes and edges...



Directed Property Graph

Nodes and relationships can have labels.

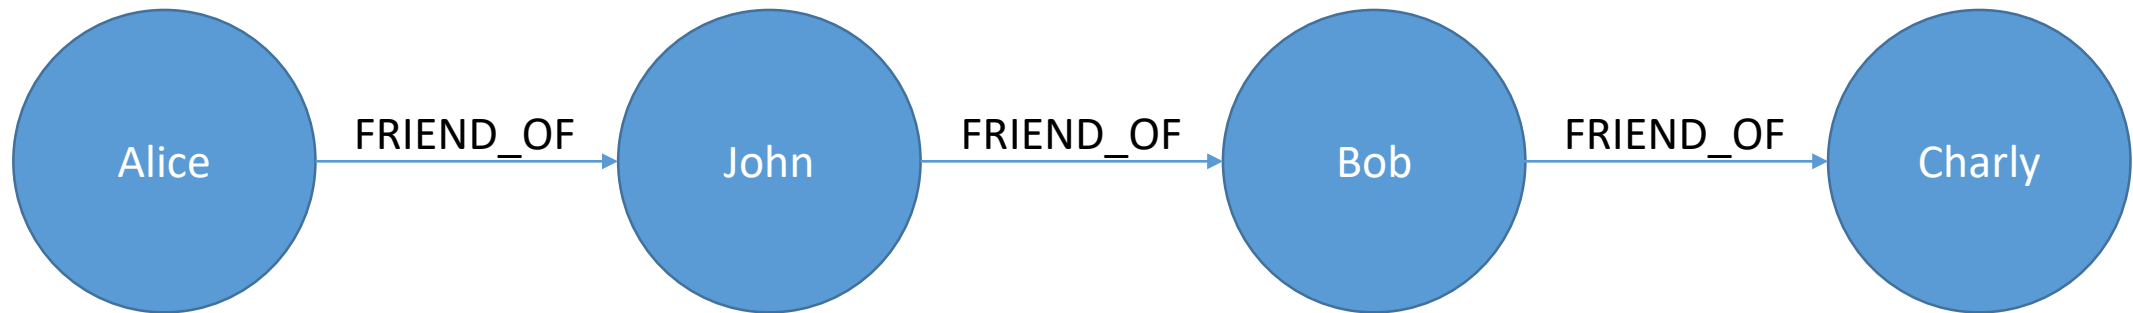
Nodes and relationships can have properties.



Graphs are schemaless (NoSQL)

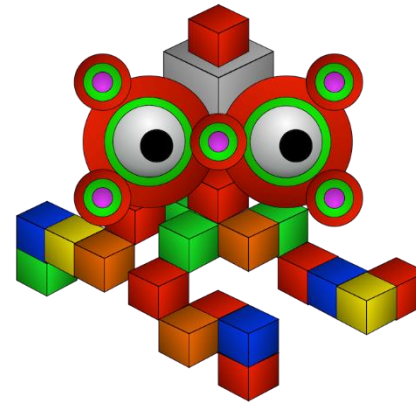
It's all about the paths

A path is a sequence of nodes connected by relations.



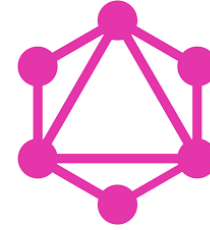
Graph databases

- **Cosmos DB** (Tinkerpop)
- **Neo4j**
- TinkerPop
- Microsoft SQL server 2017
- AllegroGraph
- OrientDB
- Many more...



Graph languages

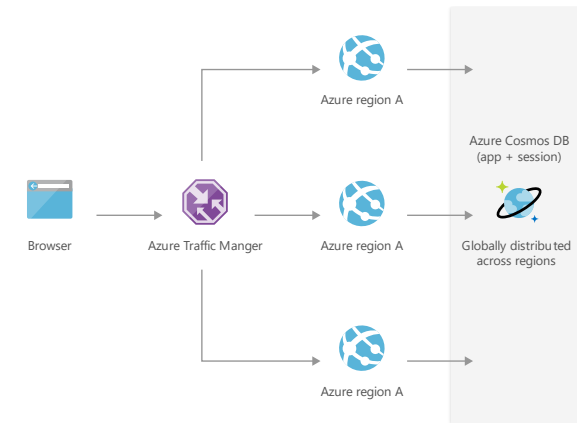
- **Gremlin** (Cosmos DB / Tinkerpop)
- **Cypher** (Neo4j)
- GraphQL
- SPARQL (Allegrograph)
- A few more...



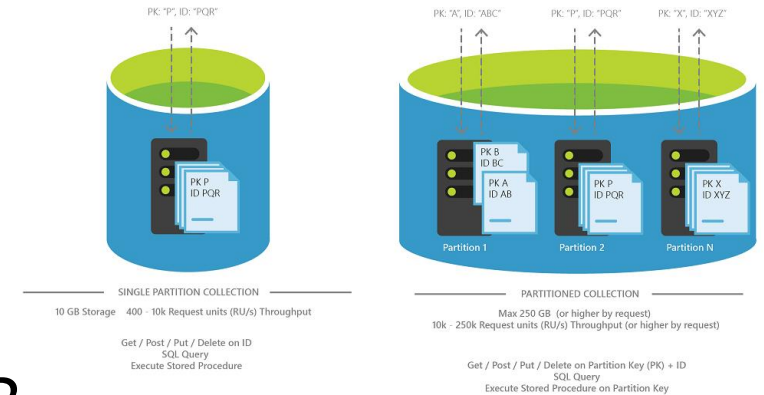
Cosmos DB / Gremlin

Microsoft's multi-model database in Azure

- Globally distributed
- Massive scale
- Guaranteed low latency
- Very high availability
- Five consistency levels
- Graph model based on Tinkerpop
- RUs / second (~ read of 1KB document)

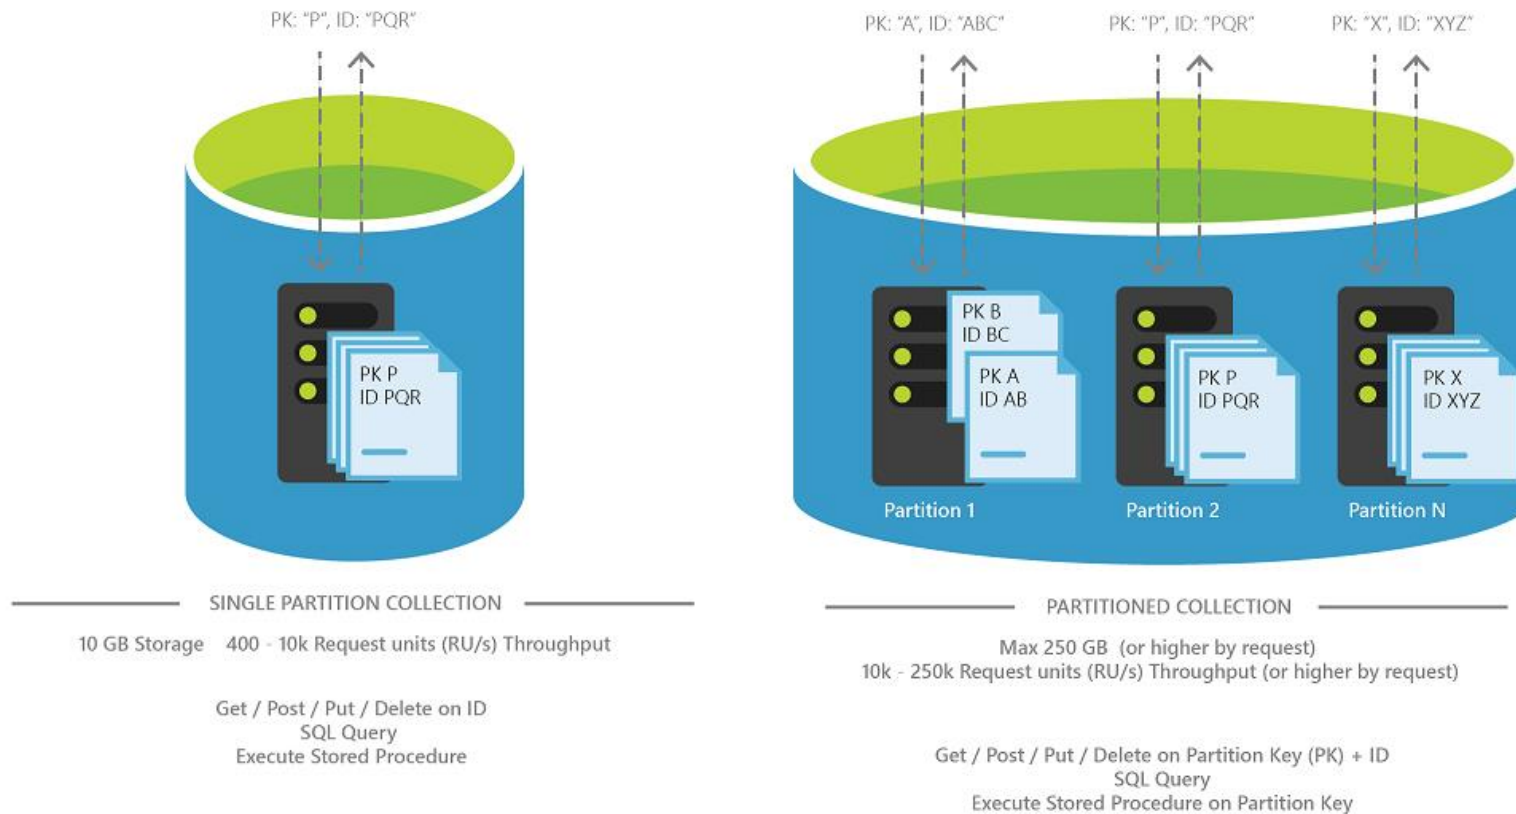


Cosmos DB - Partitioning



- Partition key optional if collection < 10GB
- Ids for Vertices and Edges must be unique per partition
- Partition key property must be present in each vertex
- Edges are stored in same partition as their out vertex
 - `john.addE('knows').to(mary)`
- PartitionKey property?
- Choose partition key to best query from a single partition?

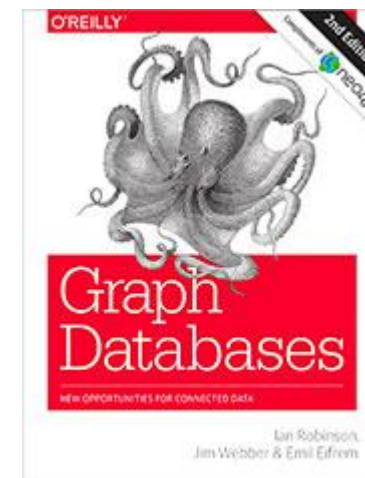
Cosmos DB - Partitioning





Neo4j / Cypher

- (One of) the oldest players in the field
- Native graph storage and processing
- Huge community
- Easy to learn
- Free (e-)book 😊
- Enterprise grade graph database
- Java, C#, Python, Javascript, Ruby, and more



What about the competition?

Why not use SQL?

- Joining tables for relation can become slow and cumbersome.

Or another NoSQL solution?

- In general, they don't support relations at all.

They may or may not be a better option for some use cases but not for...

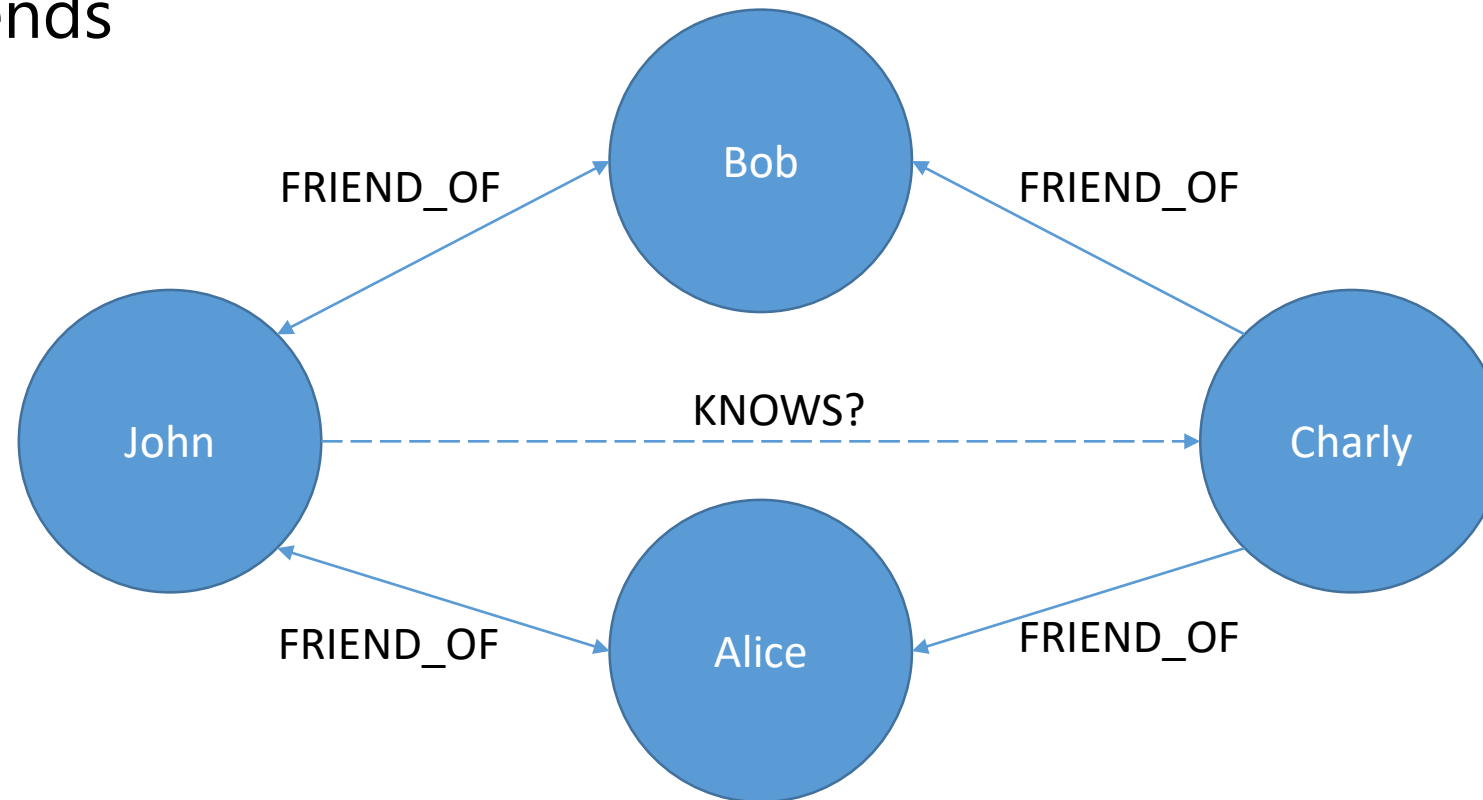
Use Cases

Use cases for Graph databases

- Social data
- Recommendations
- Fraud detection
- Geospatial
- Authorization
- Analytics

Social data

Common friends
Influencers



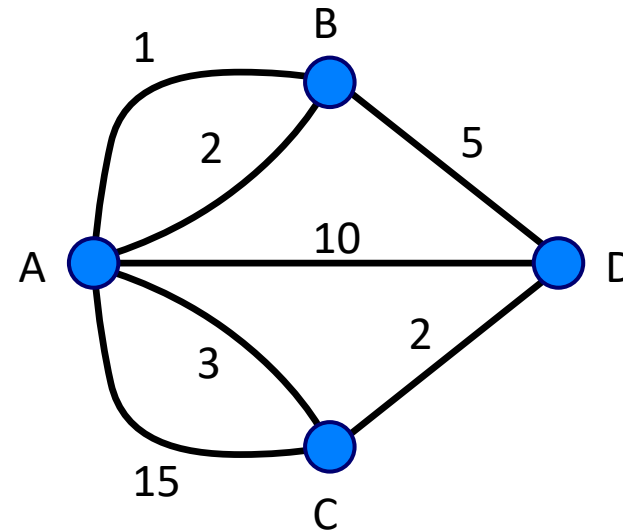
Six degrees of separation?

Geospacial

Eulers problem

Route calculations

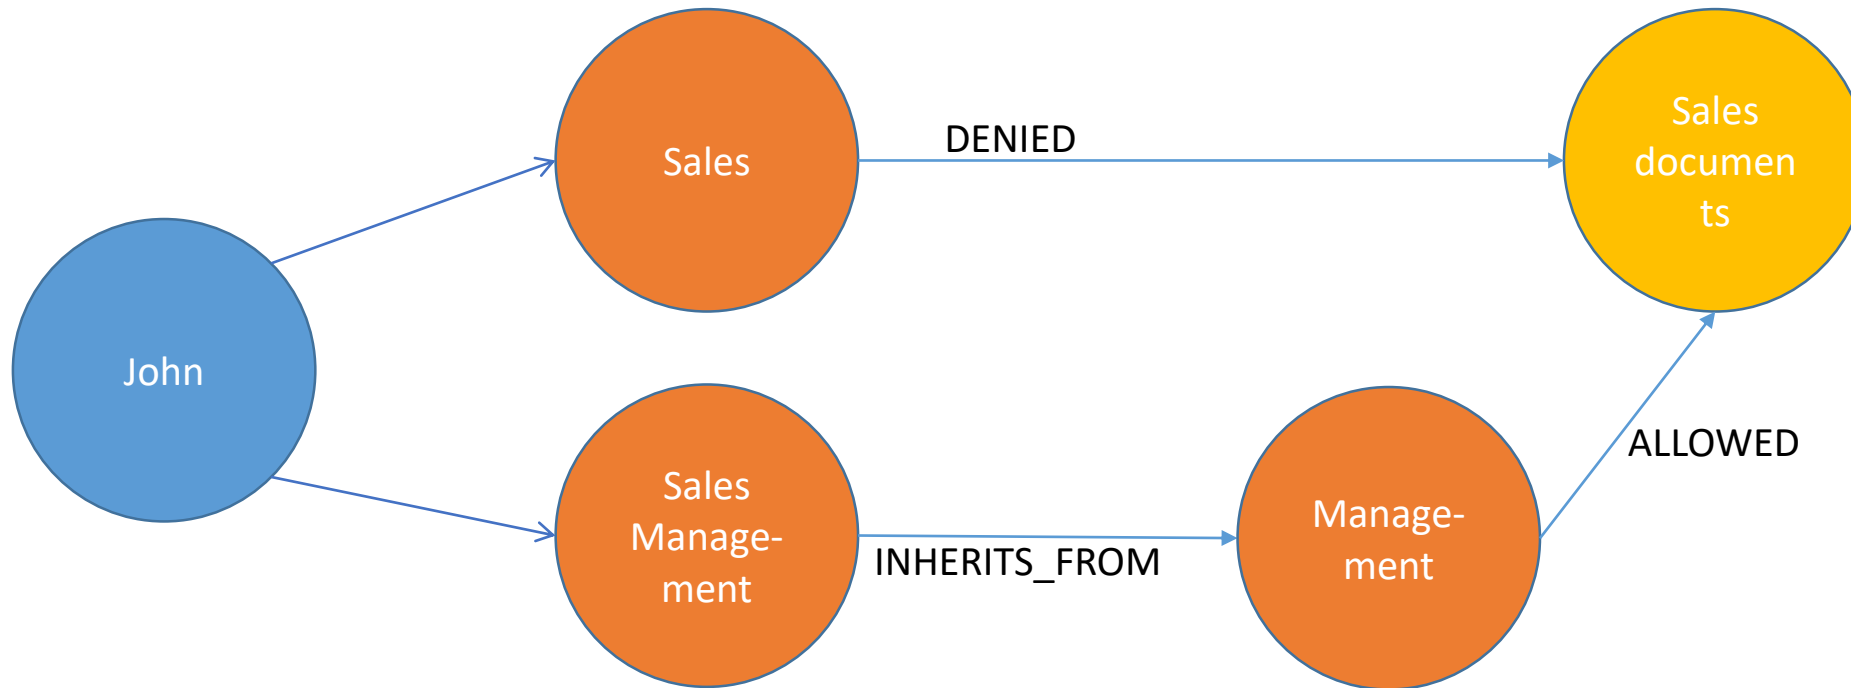
- Shortest
- Cheapest



```
MATCH p = (a:Place {name: "A"})-[:CONNECTS_TO*]->(n:Place {name: "D"})
RETURN p AS shortestPath, reduce(cost=0, r in relationships(p) | cost + r.Cost) AS totalCost
ORDER BY totalCost ASC
LIMIT 1
```

Authorization

Store and verify fine-grained access control



Tools and first code

Cosmos DB

Azure Portal

New Graph | ...

GRAPHS

- demo
 - sandbox
 - Graph
 - Scale & Settings
 - Stored Procedures
 - User Defined Functions
 - Triggers
 - Cyrtag
 - Movies

Graph

New Vertex Style

g.V() Apply Filter

Graph JSON

Results

- A
- B
- C
- D

Graph

```
graph LR; A((A)) --> B((B)); A --> C((C)); A --> D((D)); B --> D; C --> D;
```

> A

Properties

- id: A
- label: Place
- name: A

Sources

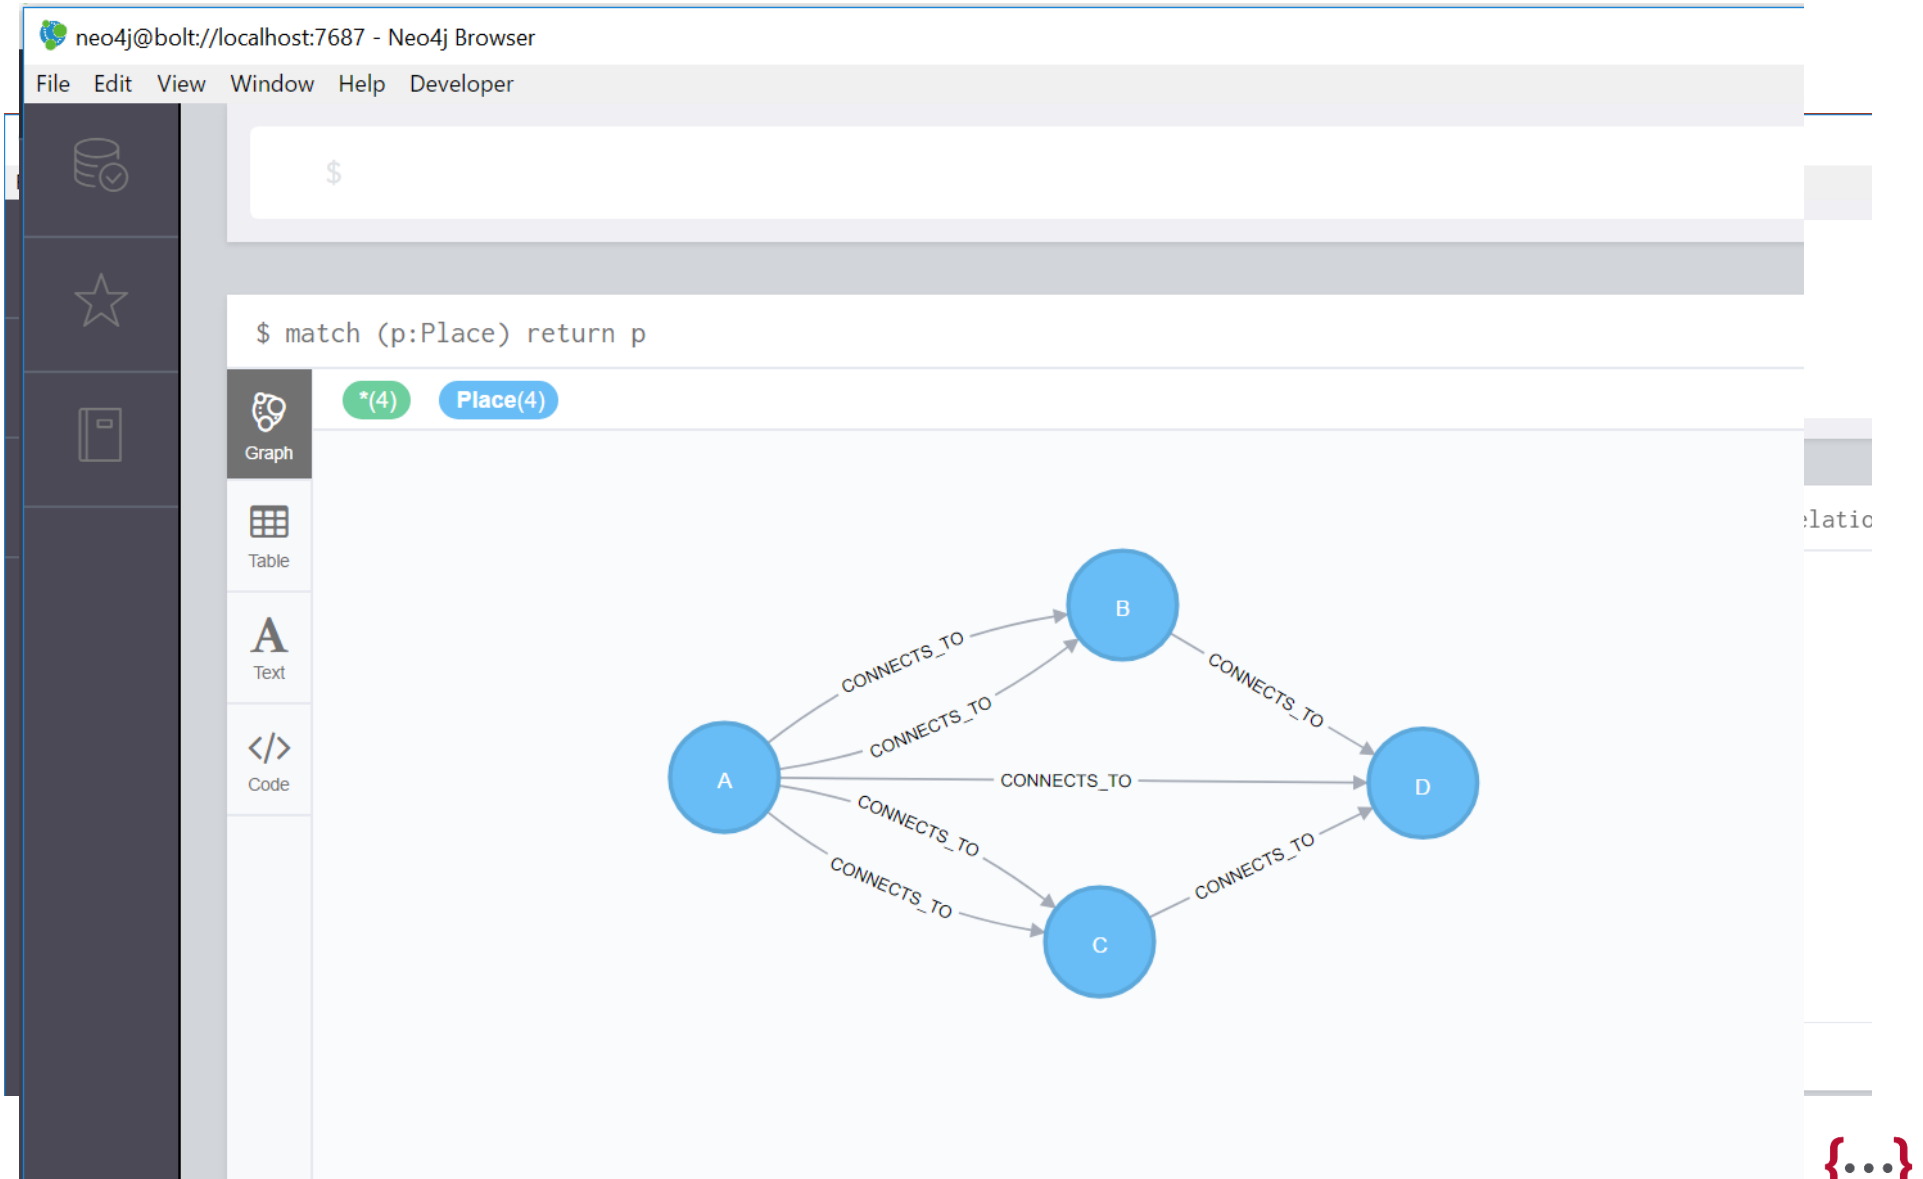
No sources found

Targets

Target	Edge label
B	CONNECTS_TO
C	CONNECTS_TO
C	CONNECTS_TO
D	CONNECTS_TO

Neo4j

Desktop Browser



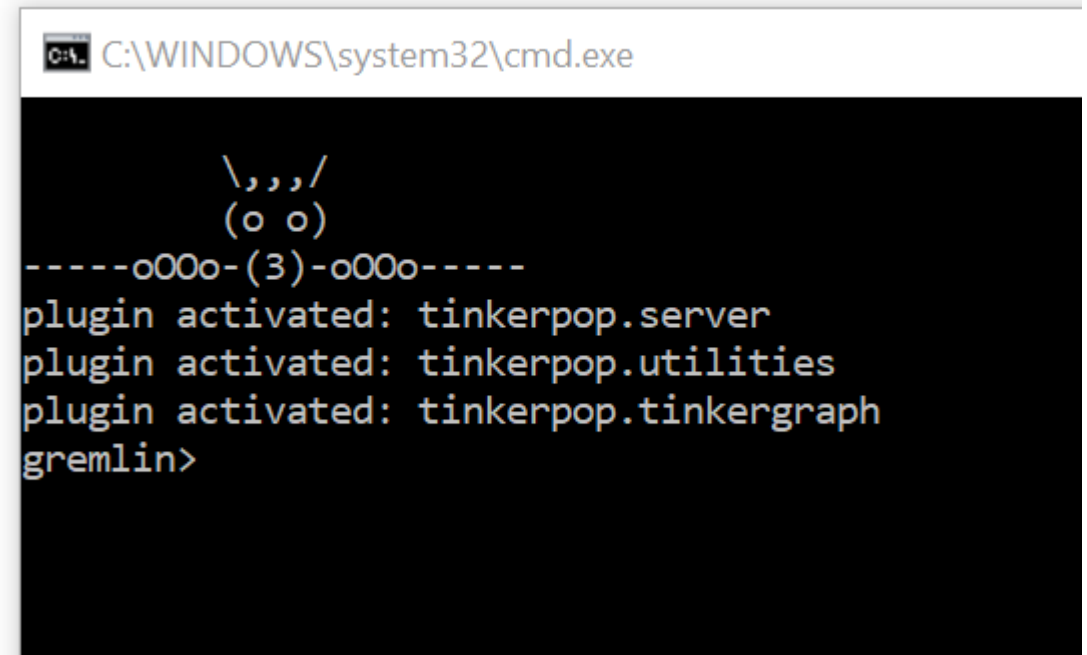
Apache Tinkerpop

Run as local Server

- Not the same as Cosmos DB!

Run local client

- Tekst based
- Can connect to Cosmos DB



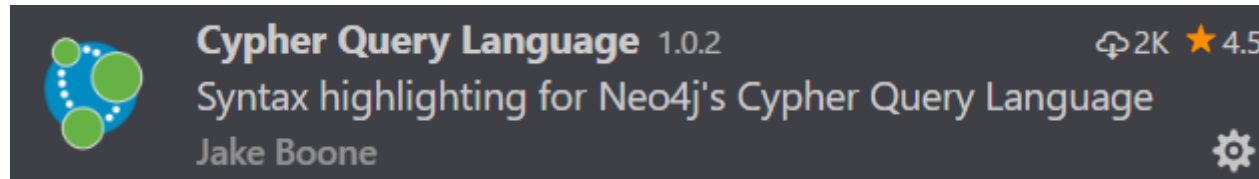
```
C:\WINDOWS\system32\cmd.exe

      \,,,\/
      (o o)
-----o00o-(3)-o00o-----
plugin activated: tinkerpop.server
plugin activated: tinkerpop.utilities
plugin activated: tinkerpop.tinkergraph
gremlin>
```

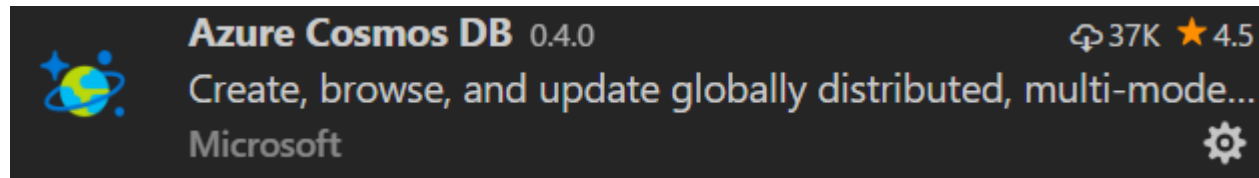

Visual Studio Code

Useful plugins:

- Cypher Query Language

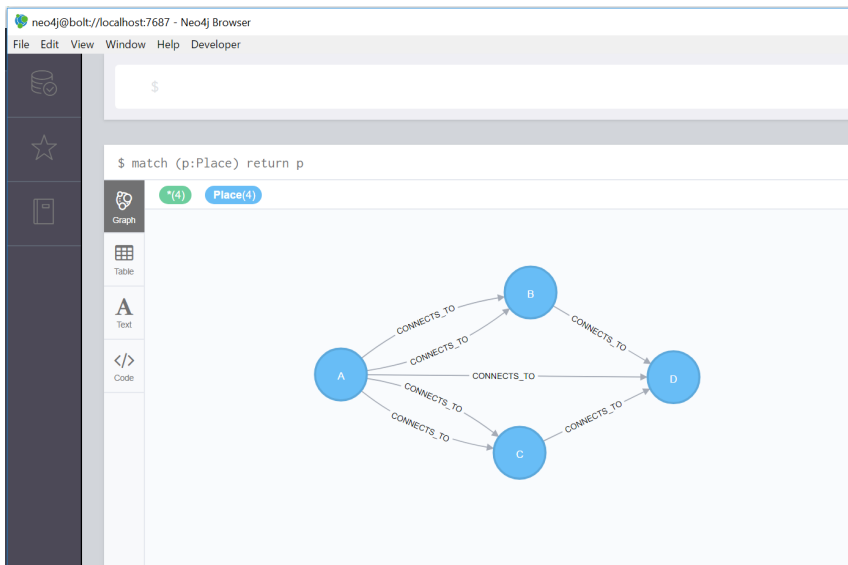


- Azure Cosmos DB



Tools demo

And a bit of
(query) code



imdb.cypher - Synwag - Visual Studio Code interface showing a Cypher query and a graph visualization. The query is: `54 USING PERIODIC COMMIT 500
55 LOAD CSV WITH HEADERS FROM "file:///title.ratings.tsv" AS line FIEL
56 WITH line
57 MATCH (t:Title {titleId: line.tconst})
58 WITH t, line
59
60
61
62
63`. The graph visualization shows a network of four nodes (A, B, C, D) connected by edges labeled "CONNECTS_TO".

```
graph LR; A((A)) -- CONNECTS_TO --> B((B)); A -- CONNECTS_TO --> C((C)); A -- CONNECTS_TO --> D((D)); B -- CONNECTS_TO --> D; C -- CONNECTS_TO --> D;
```

Code demo – Connecting to the database

```
4 references | Daniël te Winkel, 29 days ago | 1 author, 3 changes
public async Task InitialiseAsync(CancellationToken cancellationToken = default)
{
    var databaseName = _settings.Database;
    var collectionName = _settings.Collection;
    var partitionKey = _settings.PartitionKey;

    _client = new DocumentClient(new Uri(_settings.Endpoint), _settings.AuthKey);

    var database = new Database
    {
        Id = databaseName
    };
    await _client.CreateDatabaseIfNotExistsAsync(database);

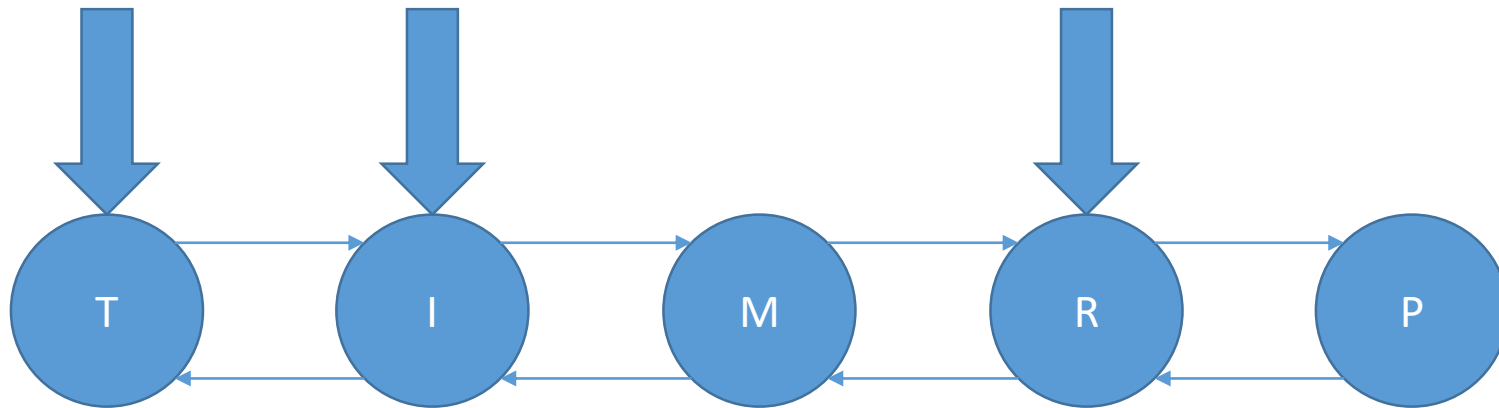
    var collection = partitionKey == default
        ? new DocumentCollection
        {
            Id = collectionName
        }
        : new DocumentCollection
        {
            Id = collectionName,
            PartitionKey = new PartitionKeyDefinition
            {
                Paths = new System.Collections.ObjectModel.Collection<string> { partitionKey }
            }
        };

    _collection = await _client.CreateDocumentCollectionIfNotExistsAsync(
        UriFactory.CreateDatabaseUri(databaseName),
        collection,
        new RequestOptions
        {
            OfferThroughput = 400,
            ConsistencyLevel = ConsistencyLevel.Session
        }
    );
}
```

Break

(Daniël)-[Needs]->(Coffee)

Some magic?



Modelling

Start with a whiteboard!

- Graphs are natural to draw; just circles and arrows.
- Use real examples
- Don't try to put too much detail in



Test in a database

Enter the examples in a database.

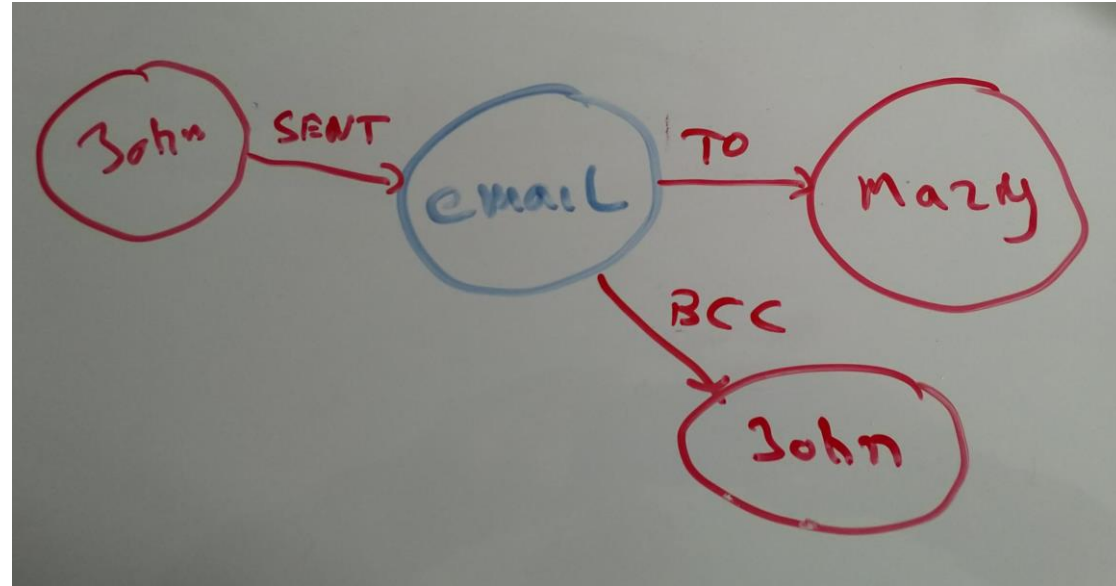
Test with queries that the business requires.

```
john = g.addV('Person').property('name', 'John')  
mary = g.addV('Person').property('name', 'Mary')  
john.addE('MAILED').to(mary)
```


Step by step

Next increment:

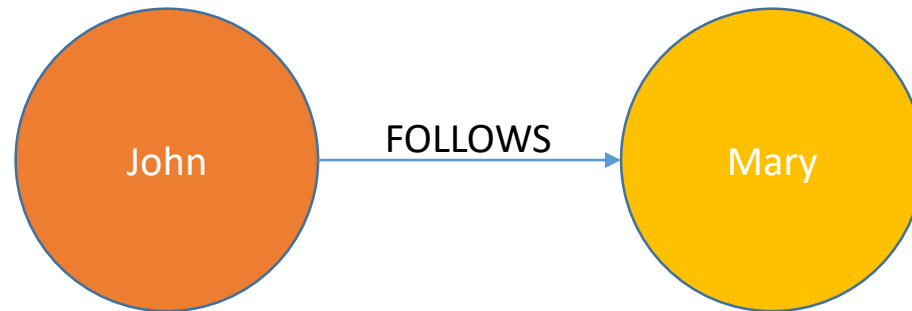
- On whiteboard
- In database
- Test queries



What is what

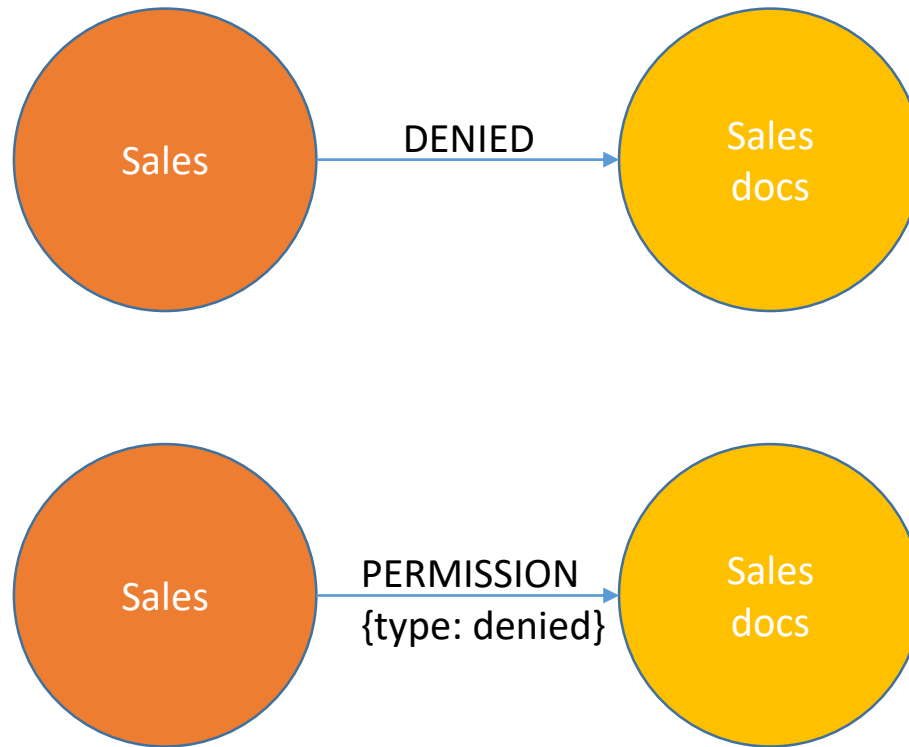
Nodes should describe **things**

Relations should describe the **relation** between the things



Labels or properties

Using specialised relationships will be faster to query



Graph languages

Gremlin

- Developed by Apache
- Functional language
- Multiple values possible for a Vertex property
- Many Dialects
 - Gremlin-Java8, Gremlin-Groovy, Gremlin-Python, e.t.c.

Cypher

- Created by Neo4j
- Open sourced as openCypher
- Graph query language
- Declarative pattern matching
- Similarities to SQL

(Neo4j)-[created]->(Cypher)

My first query

Return all vertices.

Cypher: `match (n) return n`

Gremlin: `g.V()`

Return all vertices with a given name.

Cypher: `match (n {name: "john"}) return n`

Gremlin: `g.V().has('name', 'john')`

CRUD

CRUD

- Some Gremlin
- A bit more Cypher
- And some C#

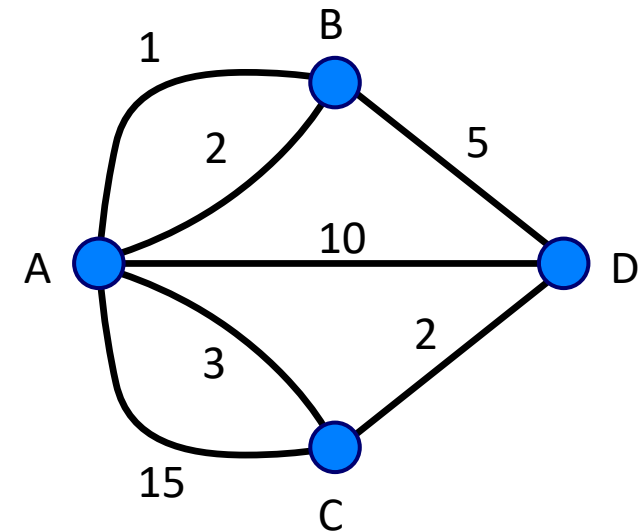
```
// Extend the model
MATCH (B:Place {name: "B"})
MATCH (C:Place {name: "C"})
MATCH (D:Place {name: "D"})
CREATE (D)-[:CONNECTS_TO {Cost: 5}]->(B)
CREATE (D)-[:CONNECTS_TO {Cost: 2}]->(C)
```

```
query = query
    .Create($"(m{pos}:Movie {{entity{pos}}})")
    .WithParam($"entity{pos}", movieEntity);
```

Querying

Querying

(Daniël) – [Shows] –> (Some queries)



The end?

Hands-on session?

If you are interested in a hands-on session to explore the (IMDB?) Graph database technology, please respond to the feedback mail



Questions?

Conclusions

Graph databases are great at querying **connected** data.

Graph databases are a great tool to have in your data toolbox!

More information

Links to the presentation and the code will be shared at:

<https://www.betabit.nl/nl/kennis/save-your-relation-with-a-graph>



Thank you for your time
and attention