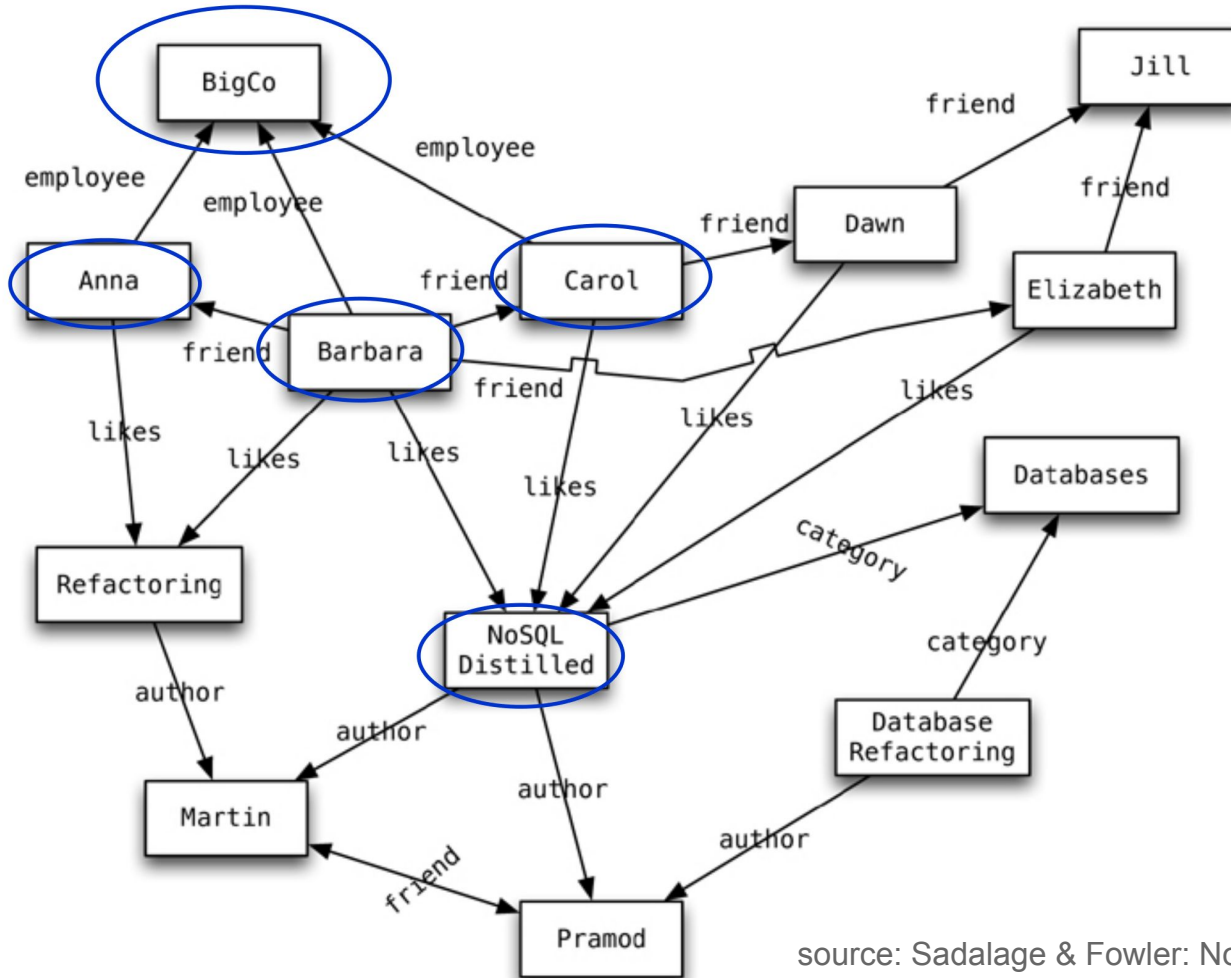
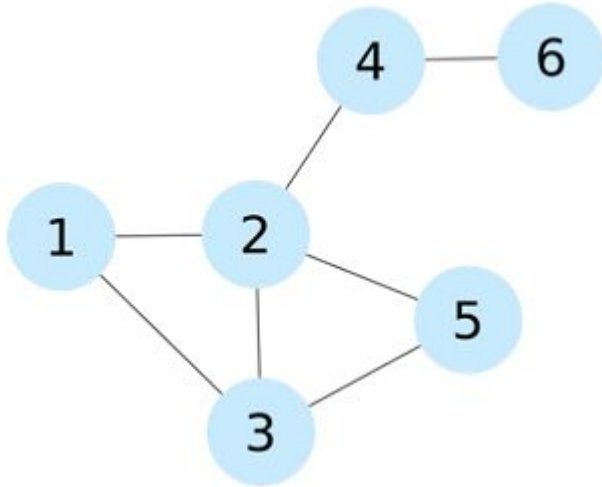


Graph Databases: Example



Adjacency Matrix: Properties



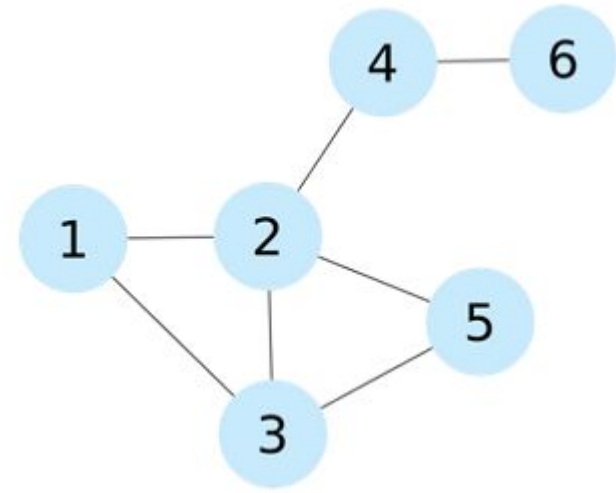
$$\begin{pmatrix} 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 1 & 0 \\ 1 & 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}$$

- Pros:
 - Adding/removing **edges**
 - **Checking** if 2 nodes are connected
- Cons:
 - Quadratic **space**: $O(n^2)$
 - We usually have **sparse** graphs
 - **Adding nodes** is expensive
 - Retrieval of **all** the **neighbouring nodes** takes linear time: $O(n)$

Data Structure: Adjacency List

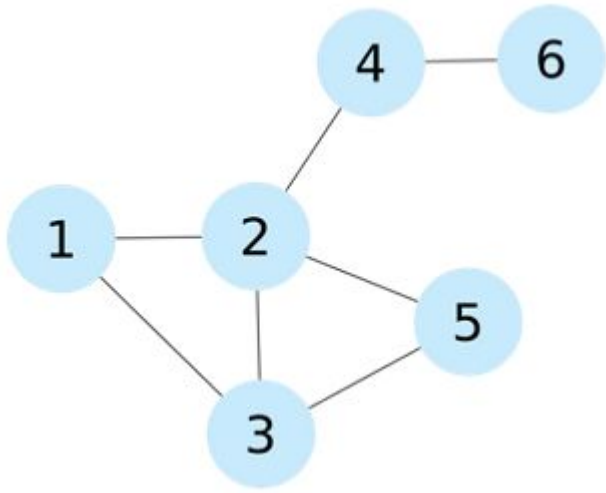


- A **set of lists**, each enumerating **neighbours** of one **node**
 - Vector of n pointers to adjacency lists
- **Undirected** graph:
 - An edge connects nodes i and j
 - \Rightarrow the adjacency list of i contains node j and **vice versa**
- Often **compressed**
 - Exploiting **regularities** in graphs



- $N1 \rightarrow \{N2, N3\}$
- $N2 \rightarrow \{N1, N3, N5\}$
- $N3 \rightarrow \{N1, N2, N5\}$
- $N4 \rightarrow \{N2, N6\}$
- $N5 \rightarrow \{N2, N3\}$
- $N6 \rightarrow \{N4\}$

Incidence Matrix: Properties



$$\begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

- Pros:
 - Representation of **hypergraphs**
 - where one **edge** connects an **arbitrary** number of nodes
- Cons:
 - Requires $n \times m$ bits (for most graphs $m \gg n$)
 - Listing neighborhood is slow

Improving Data Locality



- Performance of the **read/write** operations
 - Depends also on **physical organization** of the data
 - **Objective**: Achieve the best “data locality”
- **Spatial** locality:
 - if a data **item** has been **accessed**, the **nearby** data items are likely to be **accessed** in the following computations
 - e.g., during graph traversal
- **Strategy**:
 - in graph **adjacency matrix** representation, **exchange** rows and columns to improve the disk cache hit ratio
 - Specific **methods**: BFSL, Bandwidth of a Matrix, ...

Data Locality: Example



$$\begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 \end{pmatrix}$$

$$\begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix}$$

This matrix has **better** data **locality**, more **efficient** traversal

Breadth First Search Layout (BFSL)

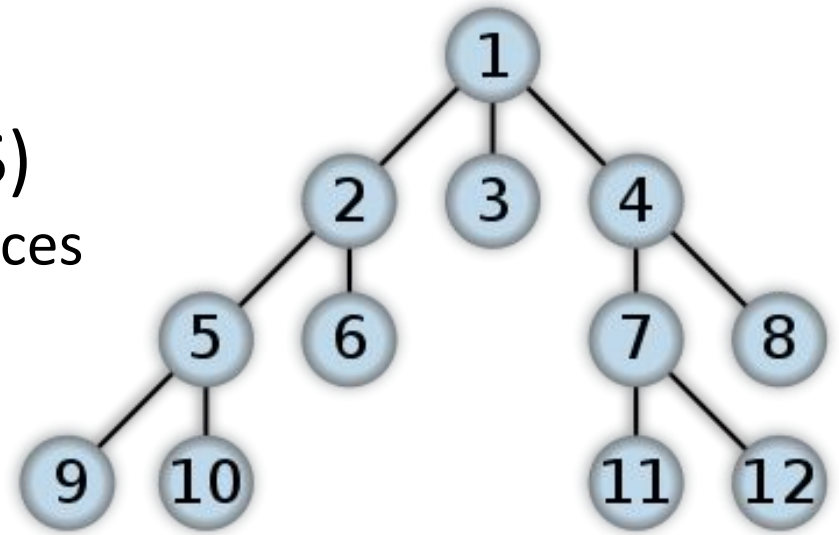


- Input: vertices of a graph
- Output: a **permutation** of the vertices
 - with better cache performance for graph traversals
- BFSL algorithm:
 1. Select a **node** (at random, the origin of the traversal)
 2. **Traverse** the graph using the BFS alg.
 - generating a list of vertex identifiers in the **order** they are **visited**
 3. Take the **generated** list as the **new** vertices **permutation**

Breadth First Search Layout (2)



- Let us recall:
Breadth First Search (BFS)
 - FIFO **queue** of **frontier** vertices

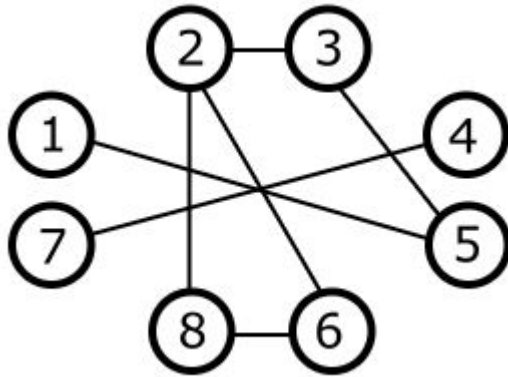


- Pros: **optimal** locality for traversal from the **root**
- Cons: starting traversal from **other nodes**
 - The further, the worse

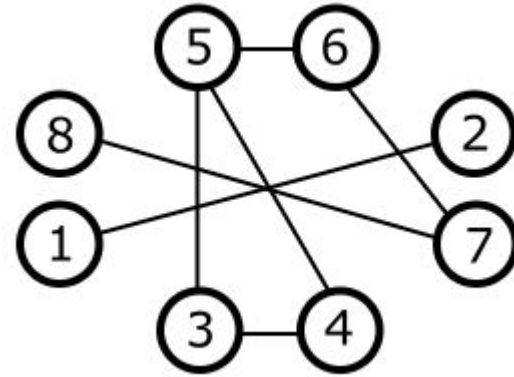
Matrix Bandwidth: Motivation



- **Graph** represented by adjacency **matrix**



$$\begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 \end{pmatrix}$$



$$\begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix}$$





A Bit of a Theory

Graph partitioning

Graph Partitioning



- Some graphs are **too large** to be fully loaded into the **main memory** of a **single** computer
 - Usage of **secondary** storage **degrades** the **performance**
 - Scalable **solution: distribute** the graph on multiple nodes
- We need to **partition** the graph reasonably
 - Usually for a particular (set of) operation(s)
 - The shortest path, finding frequent patterns, **BFS**, spanning tree search

Example: 1-Dimensional Partitioning



- Aim: **Partition** the graph to solve BFS efficiently
 - Distributed into shared-nothing parallel system
 - Partitioning of the **adjacency matrix**
- **1D partitioning of Adjacency Matrix:**
 - Matrix **rows** are randomly assigned to the P nodes (processors) in the system
 - Each **vertex** (and its **edges**) are **owned** by one processor

	1	2	3	4	5	6	7	8	9	10	11	12
1	0	0	0	0	0	0	0	0	0	1	1	0
2	0	0	1	0	0	0	0	1	0	0	0	0
3	0	1	0	0	0	0	1	1	0	0	0	0
4	0	0	0	0	1	0	0	0	0	0	0	1
5	0	0	0	1	0	0	0	0	0	0	0	1
6	0	0	0	0	0	0	1	0	0	0	1	0
7	0	0	1	0	0	1	0	1	0	1	1	0
8	0	1	1	0	0	0	1	0	0	0	0	0
9	0	0	0	0	0	0	0	0	0	0	1	1
10	1	0	0	0	0	0	1	0	0	0	1	0
11	1	0	0	0	0	1	1	0	1	1	0	0
12	0	0	0	1	1	0	0	0	1	0	0	0

Starting BFS traversal at node 1:

1. (at black) 1 -> 10, 11
visit green server
2. (at green) 10, 11 ->
 - a. 1, back to black
 - b. 6, visit red
 - c. 7,9, visit blue
 - d. 10, 11, myself
3. (at red) 6 -> 7
visit blue
3. (at blue) 7,9 ->
 - a. 3, back to black ...
 - b. 6, back to red
 - c. 8 -> 2,3, back to black
 - d. 10,11,12, back to green

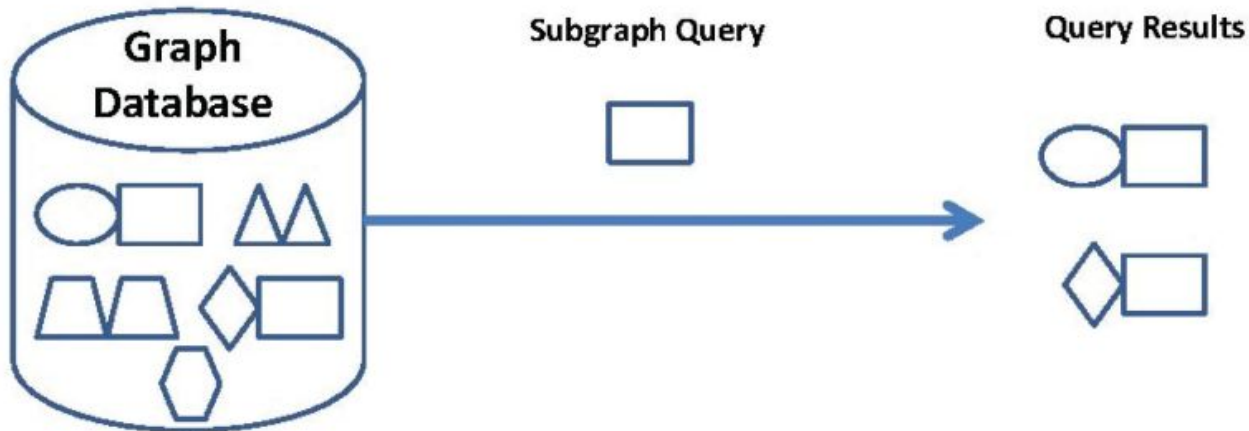
Transactional DBs: Queries



- Types of Queries

- Subgraph queries

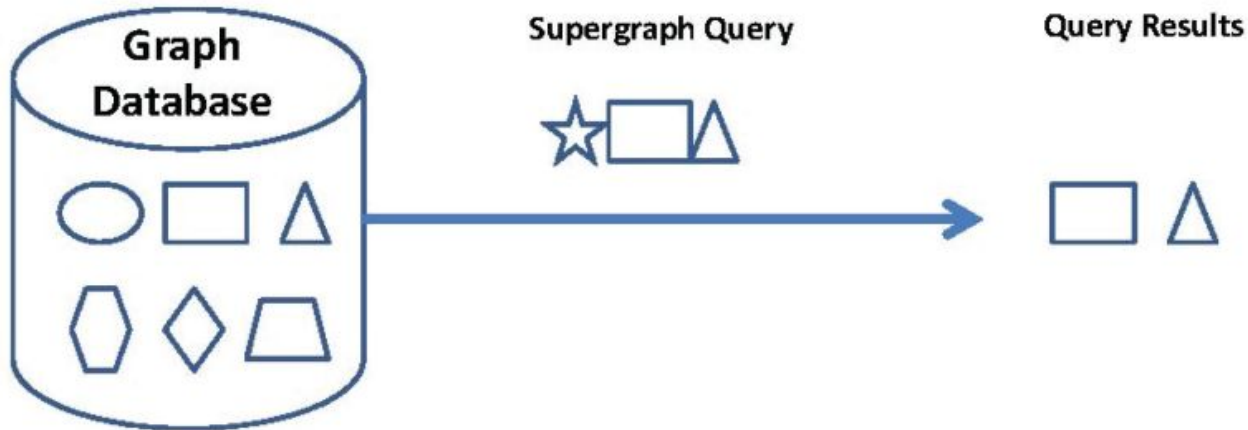
- Searches for a specific **pattern** in the graph database
 - Query = a **small graph**
 - or a graph, where some parts are uncertain, e.g., vertices with wildcard labels
 - More **general** type: allow sub-graph **isomorphism**



Transactional DBs: Queries (2)



- Super-graph queries
 - Search for graphs whose whole structure is **contained** in the **query graph**



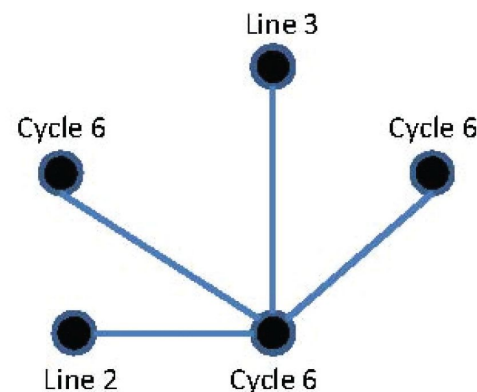
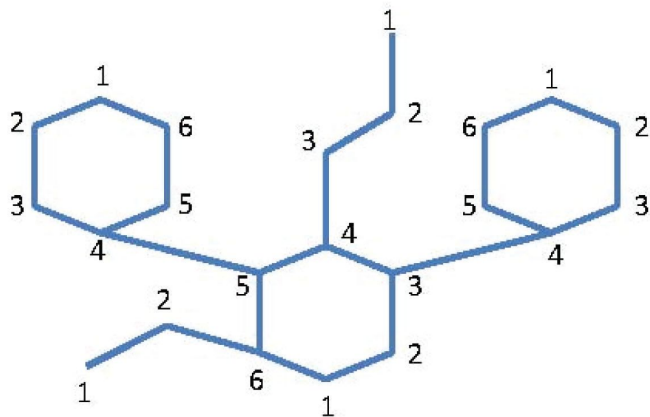
- Similarity (approximate matching) queries
 - Finds graphs which are **similar** to a given **query graph**
 - but not necessarily isomorphic
 - Key question: **how** to measure the **similarity**

Non Mining-based Techniques



- **Example: GString (2007)**

- Model the graphs in the context of organic chemistry using basic structures
 - **Line** = series of vertices connected end to end
 - **Cycle** = series of vertices that form a close loop
 - **Star** = core vertex directly connects to several vertices



Non-transactional Databases

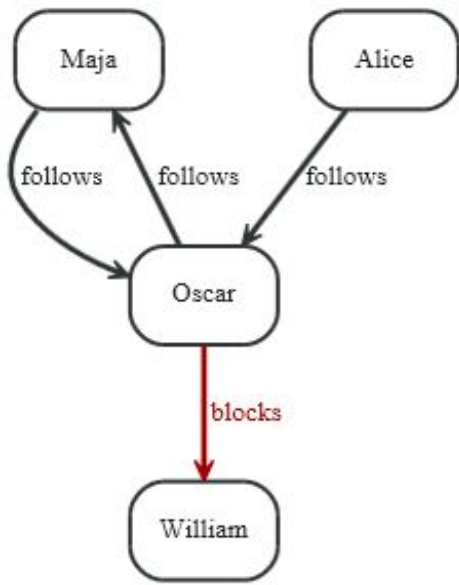


- A **few** very **large** graphs
 - e.g., Web graph, social networks, ...
- Queries:
 - Nodes/edges with properties
 - Neighboring nodes/edges
 - Paths (all, shortest, etc.)
- Our example: Neo4j

Basic Characteristics

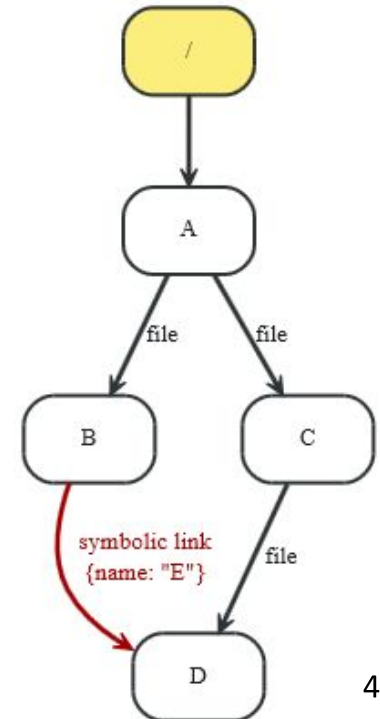


- **Different** types of **relationships** between nodes
 - To represent **relationships** between **domain** entities
 - Or to model any kind of **secondary** relationships
 - Category, path, time-trees, spatial relationships, ...
- **No limit** to the number and kind of relationships
- **Relationships** have properties
 - e.g., Since when did they become friends?



What	How
get who a person follows	outgoing <i>follows</i> relationships, depth one
get the followers of a person	incoming <i>follows</i> relationships, depth one
get who a person blocks	outgoing <i>blocks</i> relationships, depth one

What	How
get the full path of a file	incoming <i>file</i> relationships
get all paths for a file	incoming <i>file</i> and <i>symbolic link</i> relationships
get all files in a directory	outgoing <i>file</i> and <i>symbolic link</i> relationships, depth one
get all files in a directory, excluding symbolic links	outgoing <i>file</i> relationships, depth one
get all files in a directory, recursively	outgoing <i>file</i> and <i>symbolic link</i> relationships



Access to Neo4j



- **Embedded** database in Java system
- **Language**-specific connectors
 - **Libraries** to connect to a running Neo4j server
- **Cypher** query language
 - Standard language to **query** graph data
- HTTP **REST** API
- **Gremlin** graph traversal language (plugin)
- etc.

Traversal Framework



- A **traversal** is influenced by
 - Starting node(s) where the traversal begins
 - Expanders – define what to traverse
 - i.e., relationship direction and type
 - Order – depth-first / breadth-first
 - Uniqueness – visit nodes (relationships, paths) only once
 - Evaluator – what to return
 - and whether to stop or continue beyond current position

Traversal = TraversalDescription + **starting** node(s)

Traversal Framework – Java API (3)



- org.neo4j...Uniqueness
 - Indicates under what circumstances a **traversal** may **revisit** the same position in the graph

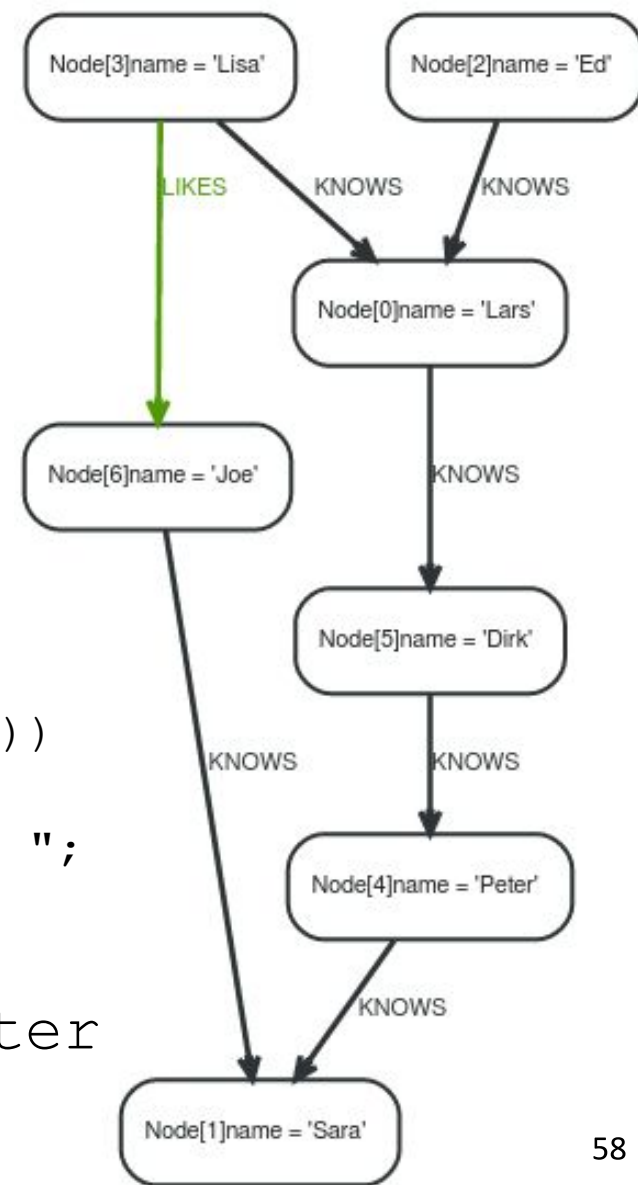
- Traverser
 - **Starts** actual **traversal** given a TraversalDescription and **starting** node(s)
 - Returns an **iterator** over “steps” in the traversal
 - Steps can be: Path (default), Node, Relationship
 - The graph is actually traversed “**lazily**” (on request)

Example of Traversal

```
TraversalDescription desc =  
    db.traversalDescription()  
        .depthFirst()  
        .relationships( Rels.KNOWS,  
                       Direction.BOTH )  
        .evaluator(Evaluators.toDepth(3));
```

```
// node is 'Ed' (Node[2])  
for (Node n : desc.traverse(node).nodes())  
{  
    output += n.getProperty("name") + ", ";  
}
```

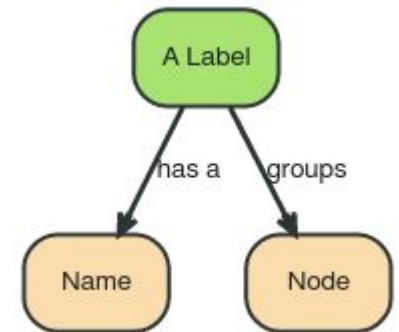
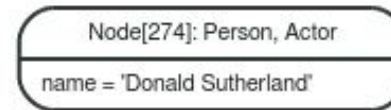
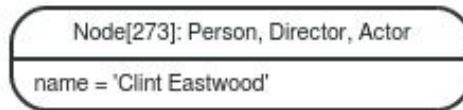
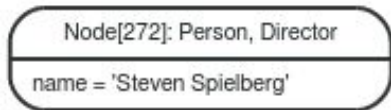
Output: Ed, Lars, Lisa, Dirk, Peter



Access to Nodes



- How to **get to the starting** node(s) before traversal
 1. Using **internal identifiers** (generated IDs)
 - **not recommended** - Neo4j generates IDs for memory objs and reuses IDs
 2. Using **properties** of nodes
 - one of the properties is typically “ID” (user-specified ID)
 - recommended, properties can be **indexed**
 - automatic indexes
 3. Using **“labels”**
 - group nodes into **“subsets”** (named graph)
 - a node can have **more** than one label
 - belong to more subsets





Neo4J: Cypher Language

Cypher Language



- Neo4j graph **query language**
 - For querying and updating
- **Declarative** – we say **what** we want
 - **Not how** to get it
 - **Not** necessary to express **traversals**
- **Human-readable**
- Inspired by SQL and SPARQL
- Still growing = syntax changes are often

Cypher: Creating Nodes (Examples)



```
CREATE (n);
```

(create a node, assign to var n)

```
Created 1 node, returned 0 rows
```

```
CREATE (a: Person {name : 'David'})
```

```
RETURN a;
```

(create a node with label 'Person' and 'name' property 'David')

```
Created 1 node, set 1 property, returned  
1 row
```

Cypher: Creating Relationships



```
MATCH (a {name:'John'}), (b {name:'Jack'})
CREATE a-[r:Friend]->b
RETURN r ;
```

(create a relation Friend between John and Jack)

Created 1 relationship, returned 1 row

```
MATCH (a {name:'John'}), (b {name:'Jack'})
CREATE a-[r:Friend {name: a.name + '->' + b.name }]->b
RETURN r
```

(set property 'name' of the relationship)

Created 1 node, set 1 property, returned 1 row



Graph Databases: When (not) to Use

Graph DBs: Suitable Use Cases



- Connected Data
 - **Social** networks
 - Any link-rich domain is well suited for graph databases
- Routing, Dispatch, and Location-Based Services
 - **Node** = **location** or address that has a delivery
 - **Graph** = **nodes** where a delivery has to be made
 - **Relationships** = **distance**
- **Recommendation** Engines
 - “your friends also bought this product”
 - “when buying this item, these others are usually bought”



Graph DBs: When Not to Use

- If we want to **update** all or a **subset** of entities
 - Changing a property on many nodes is not straightforward
 - e.g., analytics solution where all entities may need to be updated with a changed property
- **Some** graph databases may be **unable** to handle **lots** of data
 - **Distribution** of a graph is **difficult**

Questions?



References

- I. Holubová, J. Kosek, K. Minařík, D. Novák. Big Data a NoSQL databáze. Praha: Grada Publishing, 2015. 288 p.
- RNDr. Irena Holubova, Ph.D. MMF UK course NDBI040: Big Data Management and NoSQL Databases
- Sherif Sakr - Eric Pardede: Graph Data Management: Techniques and Applications
- Sadalage, P. J., & Fowler, M. (2012). NoSQL Distilled: A Brief Guide to the Emerging World of Polyglot Persistence. Addison-Wesley Professional, 192 p.
- <http://neo4j.com/docs/stable/>