

Graph DB

Venkatesh Vinayakarao

venkateshv@cmi.ac.in

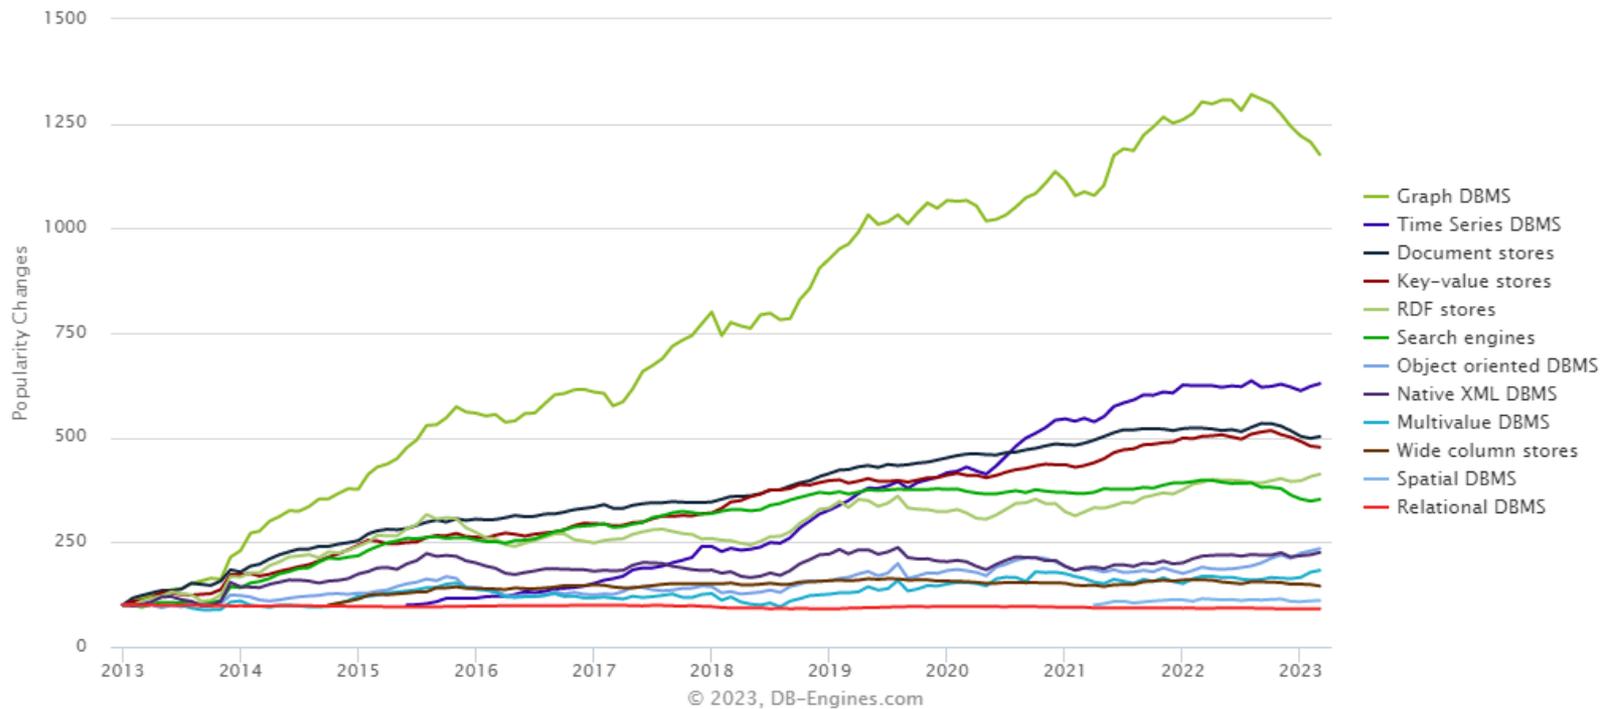
<http://vvtesh.co.in>

Chennai Mathematical Institute

We live in a connected world! . – **Neo4j**.

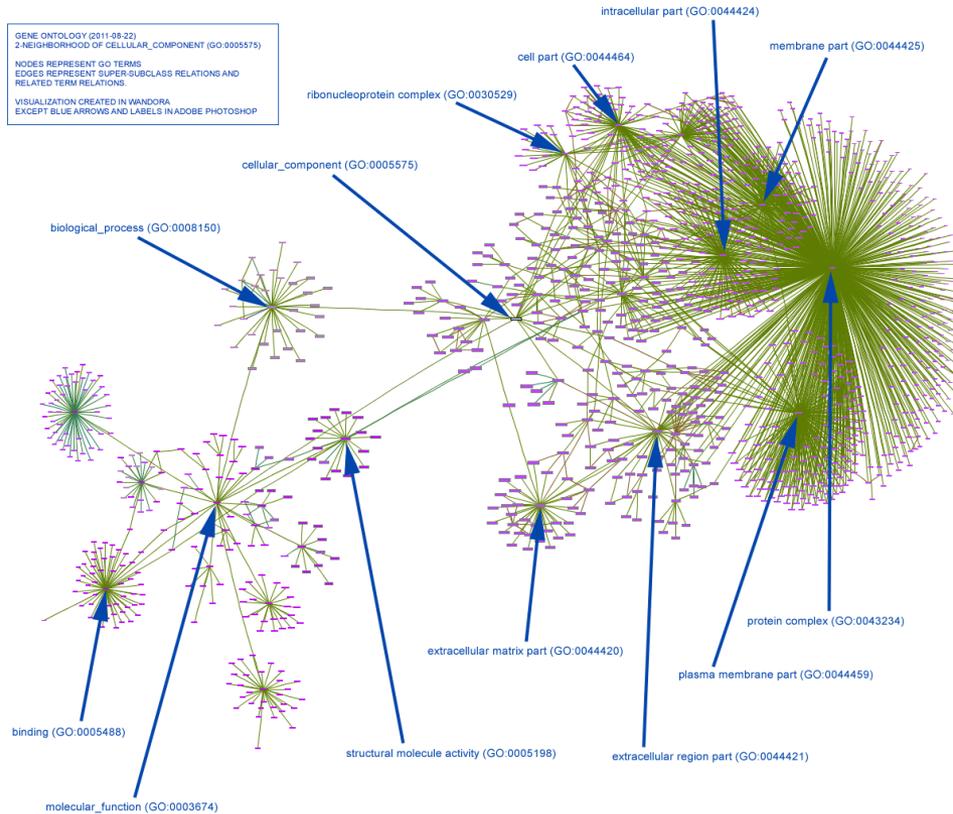
(Neo4j)-[:LOVES]-(Developers)

Change in Popularity



Source: https://db-engines.com/en/ranking_categories

Gene Ontology Model

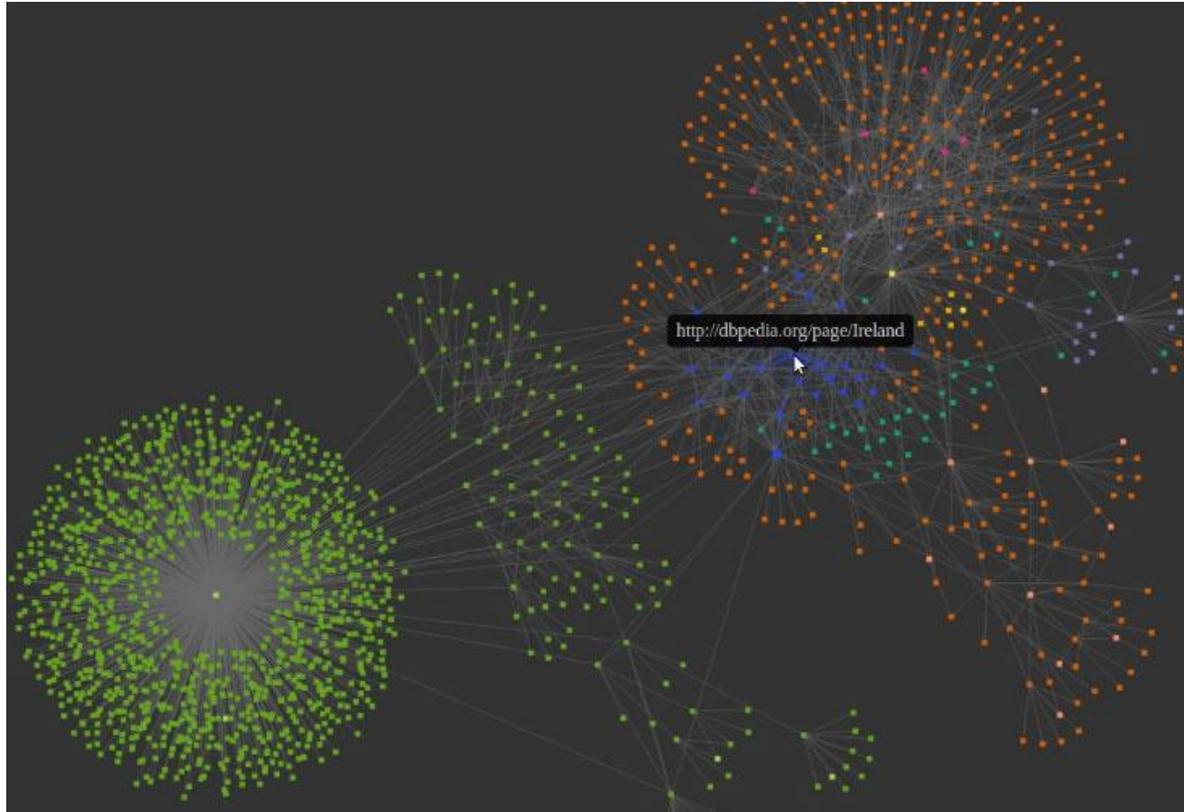


Number of topics: 177301

**Number of associations:
280198**

Source: http://wandora.org/wiki/Topic_map_conversion_of_Gene_Ontology

Knowledge Graphs



Source: <https://www.ibm.com/blogs/research/2016/01/from-knowledge-graphs-to-cognitive-computing/>

Graph Database

Relationships between data is equally as important as the data itself.

Storage and Processing

- Native Graph Storage & Processing
 - Optimized for graph related use cases such as traversals
 - Use index-free adjacency
 - Each node directly references its adjacent (neighboring) nodes
 - Does not have to move through any other type of data structures to find links between the nodes.
 - Not all graph DBs use native storage
 - Some may use other dbs such as RDBMS

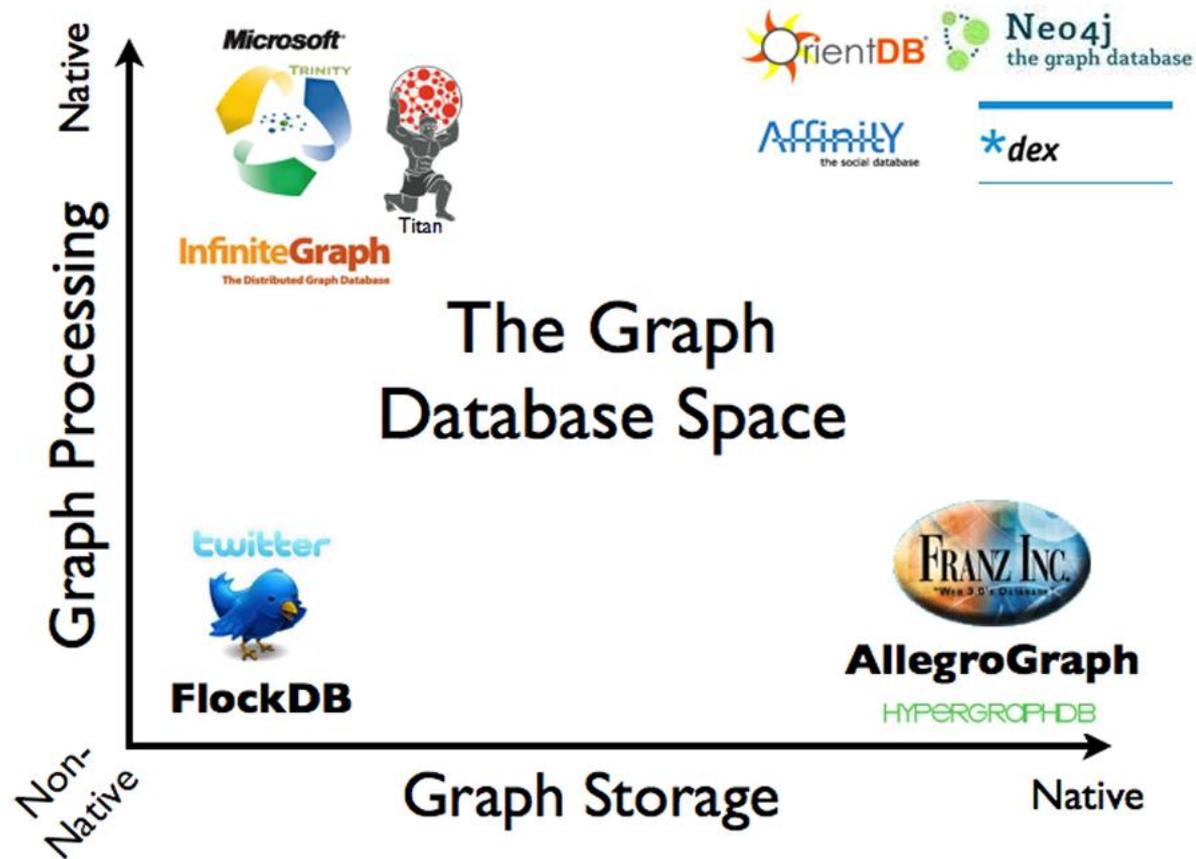
Graph Compute Engines

- Enables graph computational algorithms
 - Clustering
 - Shortest Path
 - How many flights exist between New Delhi and San Francisco having less than 3 hops?
- Distributed graph compute engines examples
 - Pegasus
 - Giraph

Why Graph DBs?

- Some data are naturally graphs
- Performance when dealing with graph data
 - Execution time for each query is proportional only to the size of the part of the graph traversed (not the size of the entire graph)
- Ease of maintenance
 - Graphs are naturally additive
 - New labels/relationships/nodes can be added without disturbing existing application features
 - Helps in agility while designing graph based applications
 - Schema free

Universe of Graph DBs



Labeled Property Graph Model

- A labelled property graph is made up of nodes, relationships, properties and labels
 - Nodes contain properties
 - Nodes can be tagged with one or more labels
 - Relationships connect nodes
 - Relationships can also have properties

Neo4j

- A leading graph database, with native graph storage and processing.
- Open Source
- NoSQL
- ACID compliant

Neo4j Sandbox

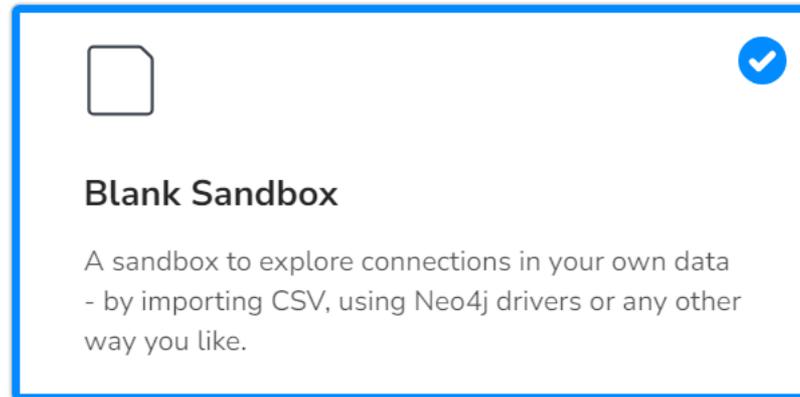
<https://sandbox.neo4j.com/>

Neo4j Desktop

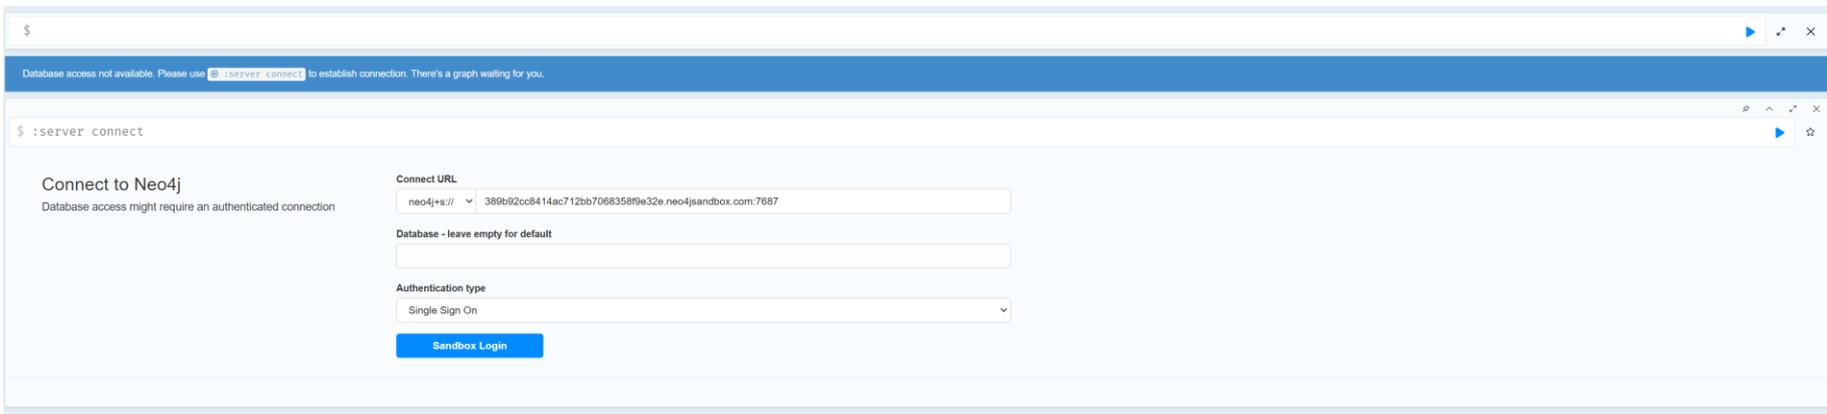
<https://neo4j.com/download>

Using Neo4J Sandbox

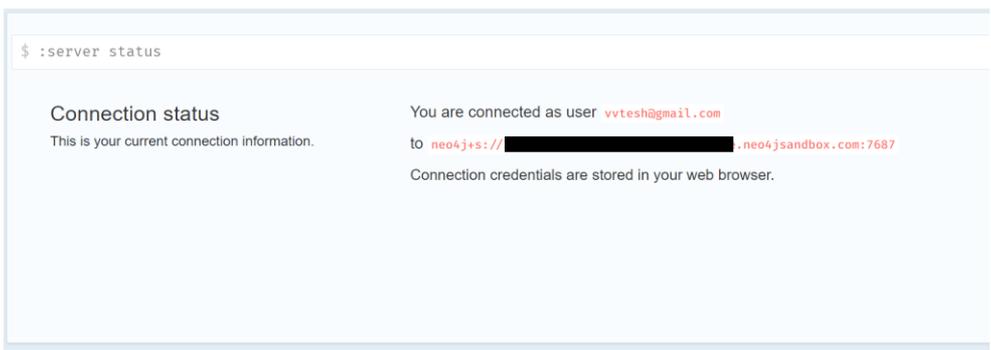
- Sign up at <https://sandbox.neo4j.com/>
- Select a project – Blank Sandbox



- Open the project



- Hit the Sandbox Login and authenticate once again



Data Model

- create (p:Person {name:'Venkatesh'})-[:Teaches]->(c:Course {name:'BigData'})

Query Language

- Cypher Query Language
 - Similar to SQL
 - Optimized for graphs
 - Used by Neo4j, SAP HANA Graph, Redis Graph, etc.

CQL

- create (p:Person {name:'Venkatesh'})-[:Teaches]->(c:Course {name:'BigData'})
- Don't forget the single quotes.

```
neo4j$ create (p:Person {name:'Venkatesh'})-[:Teaches]->(c:Course {name:'BigData'})
```



Table

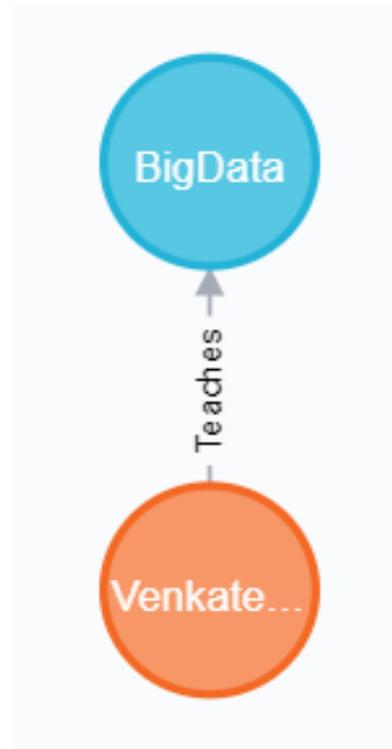


Code

Added 2 labels, created 2 nodes, set 2 properties, created 1 relationship, completed after 30 ms.

CQL

- Match (n) return n



- `match(p:Person {name:'Venkatesh'}) set p.surname='Vinayakarao' return p`

```
neo4j$ match(p:Person {name:'Venkatesh'}) set p.surname='Vinayakarao' return p
```

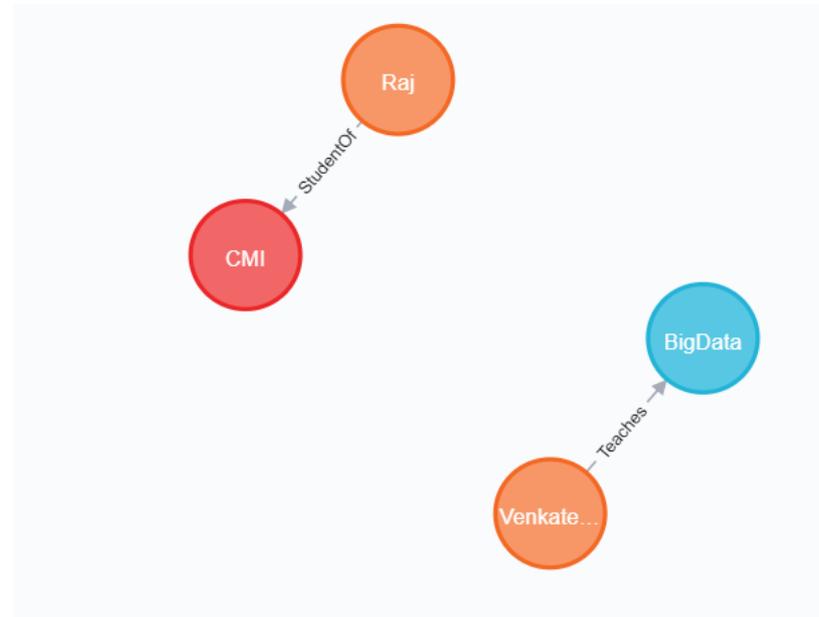


The screenshot shows a Neo4j interface with a sidebar on the left containing four icons: Graph, Table, Text, and Code. The 'Table' icon is selected and highlighted. The main area displays the variable 'p' followed by a JSON object representing the result of the query.

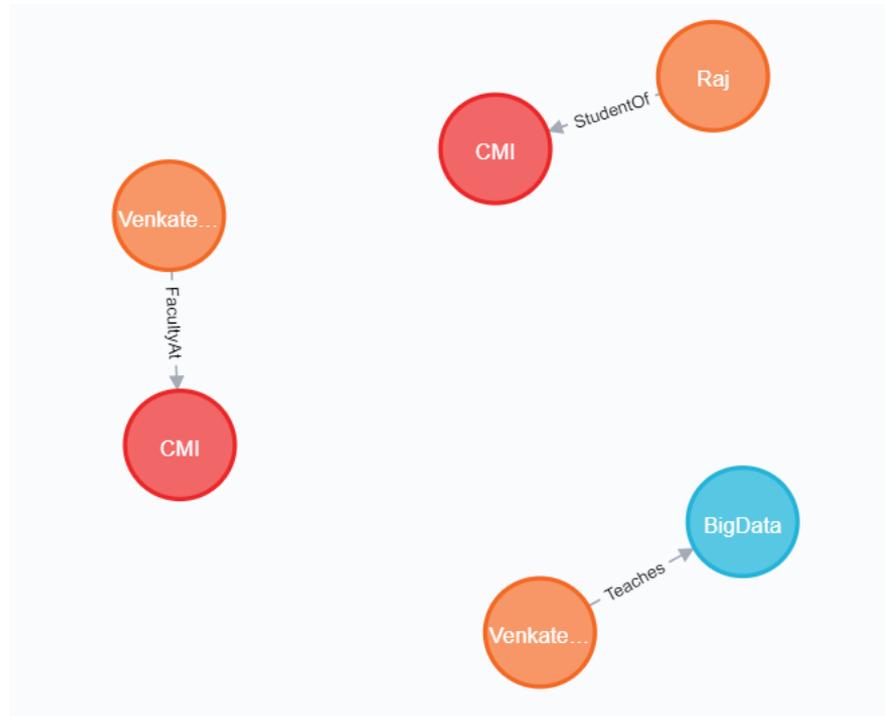
```
p
```

```
{  
  "name": "Venkatesh",  
  "surname": "Vinayakarao"  
}
```

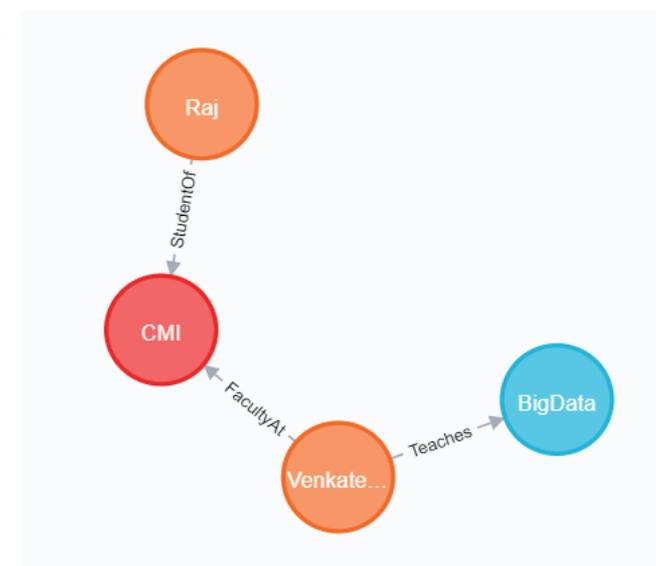
- Create (p:Person {name:'Raj'})-[StudentOf]->(o:Org {name:'CMI'})
- Match (n) return n



- create (p:Person {name:'Venkatesh'})-[:FacultyAt]->(o:Org {name:'CMI'})
- Match (n) return n



- MATCH (p:Person) where ID(p)=4
- DELETE p
- MATCH (o:Org) where ID(o)=5
- DELETE o

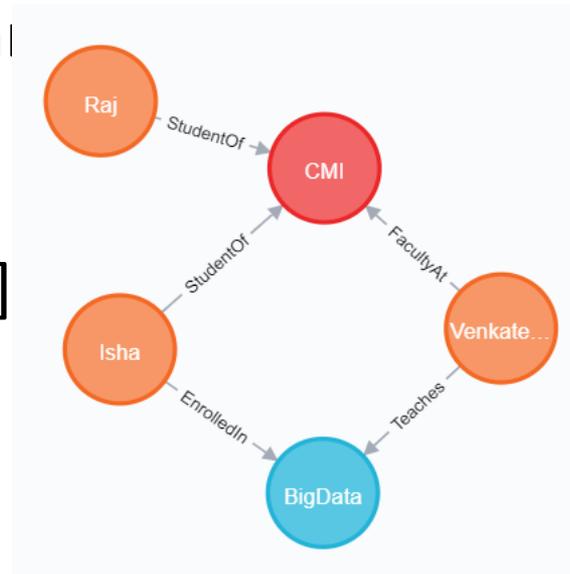


- MATCH (a:Person),(b:Org)
- WHERE a.name = 'Venkatesh' AND b.name = 'CMI'
- CREATE (a)-[:FacultyAt]->(b)

- MATCH (a:Person),(b:Course)
- WHERE a.name = 'Isha' and b.name = 'CMI'
- CREATE (a)-[:StudentOf]->(b)

- MATCH (a:Person)-[o:StudentOf] ID(o)=4
- DELETE o

- MATCH (a:Person),(b:Course)
- WHERE a.name = 'Isha' and b.name = 'BigData'
- CREATE (a)-[:EnrolledIn]->(b)



Graph Query Languages

- Just like Cypher, there are other graph query languages
 - Gremlin
 - GSQL
 - Morpheus

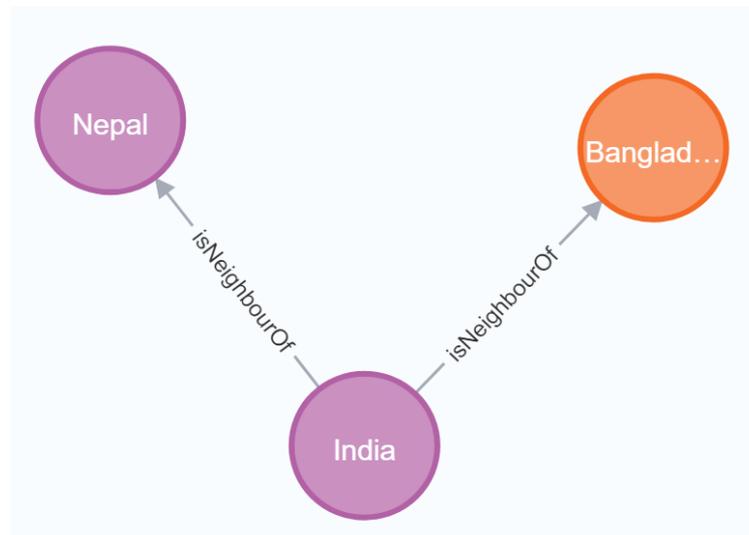
How fast are graph DBs?

- Social network: Find all friends of a user's friends, and friends of friends of friends.
- Built the query in both MySQL and Neo4j with a DB of 1M users.

Depth	Execution Time - MySQL	Execution Time - Neo4j
2	0.016	0.010
3	30.267	0.168
4	1,543.505	1.359
5	Not Finished in 1 Hour	2.132

Exercise

- Sign up at neo4j
- Create the following graph



Note

- All relationships are directional in Neo4J
- However, at query time, you may ignore the direction
 - `MATCH (x)-[:isNeighbourOf]-(y)`

Thank You