

Database Management Systems

Graph Databases

Malay Bhattacharyya

Assistant Professor

Machine Intelligence Unit
and
Centre for Artificial Intelligence and Machine Learning
Indian Statistical Institute, Kolkata

April, 2020

1 Basics

2 Property Graph Model

3 Neo4j

Basics

Graph databases use graph structures for semantic queries with nodes, edges, and properties to represent and store data.

Basics

Graph databases use graph structures for semantic queries with nodes, edges, and properties to represent and store data.

Graph databases are part of the NoSQL databases created to address the limitations of the existing relational databases.

Basics

Graph databases use graph structures for semantic queries with nodes, edges, and properties to represent and store data.

Graph databases are part of the NoSQL databases created to address the limitations of the existing relational databases.

Querying relationships within a graph database is fast because they are perpetually stored within the database itself.

Property Graph Model


Property Graph Model is a key concept in Graph Databases. In this model, data is organized as nodes, relationships, and properties (data stored on the nodes or relationships).

Property Graph Model

Property Graph Model is a key concept in Graph Databases. In this model, data is organized as nodes, relationships, and properties (data stored on the nodes or relationships).

- Nodes are the entities in the graph. They can hold any number of attributes (key-value pairs) called properties. Nodes can be tagged with labels, representing their different roles in your domain. Node labels may also serve to attach metadata (such as index or constraint information) to certain nodes.
- Relationships provide directed, named, semantically-relevant connections between a pair of node entities. A relationship always has a direction, a type, a start node, and an end node. Like nodes, relationships can also have properties. In most cases, relationships have quantitative properties (e.g., weights, costs, ratings, time intervals, etc.).

Property Graph Model – An example



Neo4j

Neo4j is an open-source, NoSQL based, native graph database that provides an ACID-compliant transactional back-end for various applications.

Neo4j efficiently implements the Property Graph Model down to the physical level (i.e., the data is stored exactly as you connect it), and the database uses pointers to navigate and traverse the graph.