

Chapter 47

Query Languages for Graph Databases

Kornelije Rabuzin
University of Zagreb, Croatia

ABSTRACT

In the past few years, many NoSQL databases have emerged, including graph databases. NoSQL databases have certain advantages and they can be used in certain domains as an alternative to relational databases. In order to use graph databases, one needs to be familiar with specific languages like Cypher Query Language (CQL) or Gremlin. However, some statements in CQL can be considered too complex for end users as it is shown later on. Because of that, the main idea of this chapter is to explore two other languages for graph databases. One of them is new and it is used to pose queries visually. Since CQL does not support recursion, views, etc., the other language is used to show how to use recursion and views on a graph database.

INTRODUCTION

One can find more or less complicated definitions, but one could say that a database is an organized collection of data. It is assumed that a database is stored as files on a computer. Actually, any list of items on a paper could be called a collection of data as well, but today it is assumed that databases are stored digitally.

In order to create a text file, some text processor should be used (for example, Word, Wordpad, Notepad, etc.). In the same way, in order to create a database, some Database Management System (DBMS) should be used. Many systems are available, including Oracle, DB2, SQL Server, PostgreSQL, MySQL, Microsoft Access, Neo4j, MongoDB, etc. The majority of DBMSs are relational in nature. “Relational” means that the main structure that is used to store data is a relation, or as users usually call it, a table. One table contains zero or more rows. The idea of storing data into *relations* (tables) is quite old and relational databases have been used for over 40 years. Relations store data in a compact and organized way and redundancy and anomalies that are caused by redundancy are usually avoided. The relational model was introduced by dr. E. F. T. Codd. in the late 1960s and early 1970s.

DOI: 10.4018/978-1-5225-7598-6.ch047

In order to use a database management system, one has to be familiar with Structured Query Language (SQL). SQL is a standardized language that is supported by all major Relational DBMS software vendors (it does not belong to any vendor in particular). SQL contains many statements to work with data. The CREATE statement is used to create different objects in the database, as well as the database itself. The INSERT, UPDATE and DELETE statements are used for data manipulation purposes. In order to retrieve data from one or more tables, SELECT statement is used. Many other, complex statements are available as well, but for our purposes, we can skip the details. Each DBMS has some (graphical) client program (one or more) that can be used to write queries (in SQL) in order to work with the database (for example, phpMyAdmin, Workbench, Navicat, pgAdmin, SQL Server Management Studio, etc.). More on SQL can be found in (Rabuzin, 2011, 2014).

Relational databases are mature technology; many DBMSs are available, SQL is standardized, etc. But in recent years things have changed. Extremely large amounts of unstructured and/or semi structured data come from different sources and relational databases can not cope with such large amounts of data in a satisfactory manner. Furthermore, certain queries require too much time to execute and sometimes such behavior is not acceptable. Because of this new types of databases are being invented (and used) on a daily basis, especially NoSQL technologies, which are becoming more and more important.

The word NoSQL had different meanings in the past (for example, No to SQL), but today people usually mean “Not Only SQL”. When one talks about NoSQL databases, several different types can be distinguished (some authors believe that XML databases should be included as well, as well as some other types, but it doesn’t make much difference):

- Document oriented databases,
- Column oriented databases,
- Key Value databases,
- Graph databases.

Document oriented databases, like MongoDB, do not use tables to store data. Instead, they use collections that store documents. Since one document contains all the data that are relevant, foreign keys are not used (at least not in a form that one is used to in relational databases). Document databases can use references to link between documents, but usually all the referenced data could/should be included within the document. So basically, collections would correspond to tables and documents would correspond to rows. It is important to have in mind that the document schema is flexible and each document within the collection can have different schema.

Column oriented databases use tables as well, but data is organized and stored differently. Unlike relational databases that store together values that belong to a single row (usually several rows are stored within a single unit of storage), column oriented databases store together values (for different rows) that belong to the same column. This is easy to justify. Namely, when one poses a query, in most cases one doesn’t need all the columns in the table. Instead, only a few columns are usually selected. In data warehouses dimension tables can have very large number of columns. So queries that are posed on column oriented databases should be faster because less data has to be read from the hard disk drive since column data is stored together (and not values that belong to a single row).

Key Value (KV) databases store values for defined keys. Values can be simple as well as complex. Key value databases may look trivial, but sometimes they can be very useful.

13 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:
www.igi-global.com/chapter/query-languages-for-graph-databases/214649

Related Content

Zero-Crossing Analysis and Information Divergence of Lévy Walks for Real-Time Feature Extraction

Jesus David Terrazas Gonzalez and Witold Kinsner (2016). *International Journal of Handheld Computing Research* (pp. 41-59).

www.irma-international.org/article/zero-crossing-analysis-and-information-divergence-of-ly-walks-for-real-time-feature-extraction/176418

Speech-Centric Multimodal User Interface Design in Mobile Technology

Dong Yu and Li Deng (2008). *Handbook of Research on User Interface Design and Evaluation for Mobile Technology* (pp. 461-477).

www.irma-international.org/chapter/speech-centric-multimodal-user-interface/21847

Machine Learning-Based Coding Decision Making in H.265/HEVC CTU Division and Intra Prediction

Wenchan Jiang, Ming Yang, Ying Xie and Zhigang Li (2020). *International Journal of Mobile Computing and Multimedia Communications* (pp. 41-60).

www.irma-international.org/article/machine-learning-based-coding-decision-making-in-h265hevc-ctu-division-and-intra-prediction/255093

Context-Aware Mobile Capture and Sharing of Video Clips

Janne Lahti, Utz Westermann, Marko Palola and Johannes Peltola (2009). *Mobile Computing: Concepts, Methodologies, Tools, and Applications* (pp. 1080-1095).

www.irma-international.org/chapter/context-aware-mobile-capture-sharing/26571

A Location-Tracking Method With a Convolutional Neural Network

Shiori Kawakami, Shinji Sakamoto and Shusuke Okamoto (2021). *International Journal of Mobile Computing and Multimedia Communications* (pp. 17-26).

www.irma-international.org/article/a-location-tracking-method-with-a-convolutional-neural-network/284391