MongoDB as a Graph Database

Marvin Arroz, Andersen Chen, Jonathan Lessinger, Debjani Mitra

Abstract

MongoDB, with its flexible BSON document storage, covers the representation of normalized relational data forced by RDBMSs, but also extends to cover that data denormalized to match expected application use and mapped to programming language objects. That denormalization, matching, and mapping can result in programmer productivity and enormous performance improvement that opens up new possibilities for programming languages applied to big data. In this project, MongoDB was applied to graph databases to achieve programmer productivity and performance gains. A framework was written in javascript to allow MongoDB to emulate a graph data store. Additionally, a web application was provided to allow a user to benchmark the performance of the MongoDB graph database implementation against Neo4j, an open-source graph database. In an example graph dataset with standard graph search algorithms, queries in the MongoDB framework ran significantly faster than their Neo4j counterparts. This demonstrates that MongoDB is an effective choice for handling graph data and applications.