

The rise of Big Data has highlighted the limitations of relational databases while handling large datasets, leading to the growth of NoSQL databases. This has made DBMS benchmarking crucial for performance evaluation and decision-making.

This thesis compares relational (MySQL, SQLite), graph (Neo4j, ArangoDB), document (MongoDB), and column-family (Cassandra) databases. We analyze the expressive power of their query languages and their runtime efficiency across varying data sizes. We conclude, that there's no "number one" solution for all use cases. The choice depends on factors like data volume, query complexity, and the need for joins.

For complex queries and frequent joins, MySQL and SQLite are the most expressive but may struggle with very large datasets. Cassandra and MongoDB excel in performance and scalability but require efficient schema design and targeted data redundancy. ArangoDB presents a versatile option capable of handling multiple data models but might require further investigation into its performance compared to Neo4j.