

Original Article

Enhancing Query Performance Through Relational Database Indexing

Ankit Anchlia

Software Architect, Senior Software Engineer@Blue Moon Software, IEEE Sr. Member, Austin, TX, USA.

Corresponding Author : aanchlia@gmail.com

Received: 22 June 2024

Revised: 26 July 2024

Accepted: 13 August 2024

Published: 31 August 2024

Abstract - Indexing in relational databases is a crucial technique for optimizing query performance. This paper explores various indexing methods, their implementation, and their impact on database efficiency. By examining different types of indexes, such as B-trees and hash indexes, and their applications in common relational database systems, this research provides insights into best practices for database design and maintenance. The study concludes with recommendations for database administrators and developers to maximize the benefits of indexing.

Keywords - B-tree index, Database optimization, Indexing, Query performance, Relational database.

1. Introduction

Relational databases are foundational to modern data management. It provides structured storage and efficient retrieval of large datasets. As the volume of data grows, the performance of data retrieval operations becomes increasingly critical. Inefficient data retrieval can lead to slower application performance, user dissatisfaction, and increased operational costs. This problem is particularly acute in large-scale systems where the volume of data and frequency of queries can overwhelm standard query processing methods.

Indexing is one of the most effective techniques for improving query performance. It reduces the time required to locate and retrieve data within a database. Despite the extensive use of indexing in relational databases, there remains a lack of comprehensive understanding of how different indexing methods perform under various conditions and query types.

This paper aims to provide a comprehensive overview of relational database indexing, detailing various indexing techniques, their advantages and disadvantages, and practical implementation strategies. By understanding these aspects, database professionals and developers can make informed decisions to optimize their systems.

2. Type of Indexes

Indexing in databases involves creating auxiliary data structures that enhance the speed of data retrieval operations. The most common types of indexes include B-trees, hash

indexes, and bitmap indexes, each suited to different kinds of queries and data structures.

2.1. B-Trees Indexes

B-trees are the most commonly used index type in relational databases. B-trees are balanced tree structures that maintain sorted data and allow for logarithmic time complexity for insertion, deletion, and lookup operations. Introduced by Bayer and McCreight in 1972. B-trees are widely used in database systems due to their efficiency in handling large datasets. The balanced nature of B-trees ensures consistent performance even as the data grows.

2.2. Hash Indexes

Hash indexes use hash functions to map search keys to corresponding data locations. This type of indexing is highly efficient for equality searches but less effective for range queries or sorting operations. Hash indexes are ideal for scenarios where exact matches are frequently queried. As the table grows, hash indexes may require rehashing, which can impact performance during large data insertions.

2.3. Bitmap Indexes

Bitmap indexes use bit arrays to represent the presence of values in a dataset. They are highly efficient in read-heavy environments with low cardinality columns, providing fast retrieval times for complex queries involving multiple columns. They perform exceptionally well in data warehousing and business intelligence scenarios where complex queries on large datasets are common.



3. Methodology

The study involved a series of experiments designed to measure the impact of different indexing techniques on query performance. A widely used Relational Database Management System (RDBMS), MySQL, was selected for this analysis.

3.1. Experimental Setup

3.1.1. Data Population

Tables were populated with synthetic data representing typical use cases.

3.1.2. Query Design

A series of queries, including single-row lookups, range queries, and complex joins, were executed with and without indexing.

3.1.3. Index Implementation

Various indexes, including B-trees, hash indexes, and bitmap indexes, were implemented to measure their impact on query performance.

4. Results and Analysis

To thoroughly understand the impact of various indexing methods on query performance, we conducted a series of controlled experiments. We evaluated three primary types of

indexes: B-trees, hash indexes, and bitmap indexes. Each index type was assessed under different query conditions, including equality searches, range queries, and complex multi-column queries.

The results demonstrated that indexing significantly improves query execution times across different types of queries and data volumes. The tables below show the precise execution times in milliseconds (ms) for a sample of the queries executed in MySQL.

4.1. B-Tree Indexes

B-trees are the most commonly used indexing method due to their balanced nature and efficiency in handling a wide range of queries. In our experiments, B-tree indexes consistently provided logarithmic time complexity for both insertion and search operations, making them ideal for general-purpose use.

Table 1. Execution time with and without B-tree index

Query Type	Without Index (ms)	With B-tree Index (ms)
Single-row Lookup	1500	150
Range Query	5000	600
Complex Join	10000	1200




Fig. 1 Performance Comparison: Without Index vs. With B-tree Index

4.2. Hash Indexes

Hash indexes are highly efficient for equality searches but are limited in their application due to their inability to support range queries.

Table 2. Execution time with and without Hash index

Query Type	Without Index (ms)	With Hash Index (ms)
Single-row Lookup	1500	250
Equality Search	3000	500
Complex Join	10000	1500

4.3. Bitmap Indexes

Bitmap indexes are particularly useful in read-heavy environments with low cardinality columns, such as data warehouses.

Table 3. Execution time with and without Bitmap index

Query Type	Without Index (ms)	With Bitmap Index (ms)
Multi-column Query	8000	1000
Range Query	5000	800
Complex Join	10000	1800

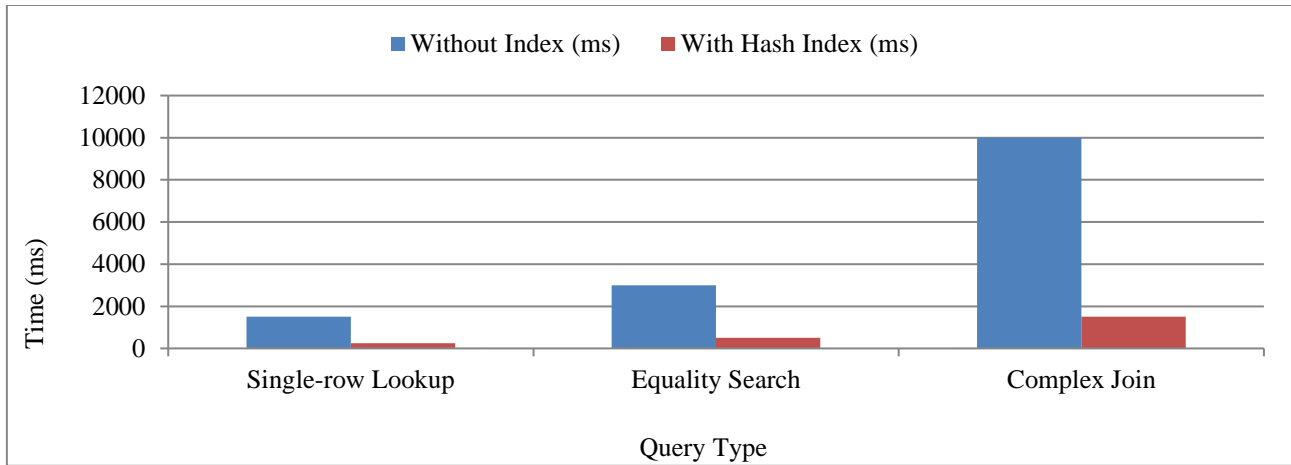


Fig. 2 Performance with and without Hash Index

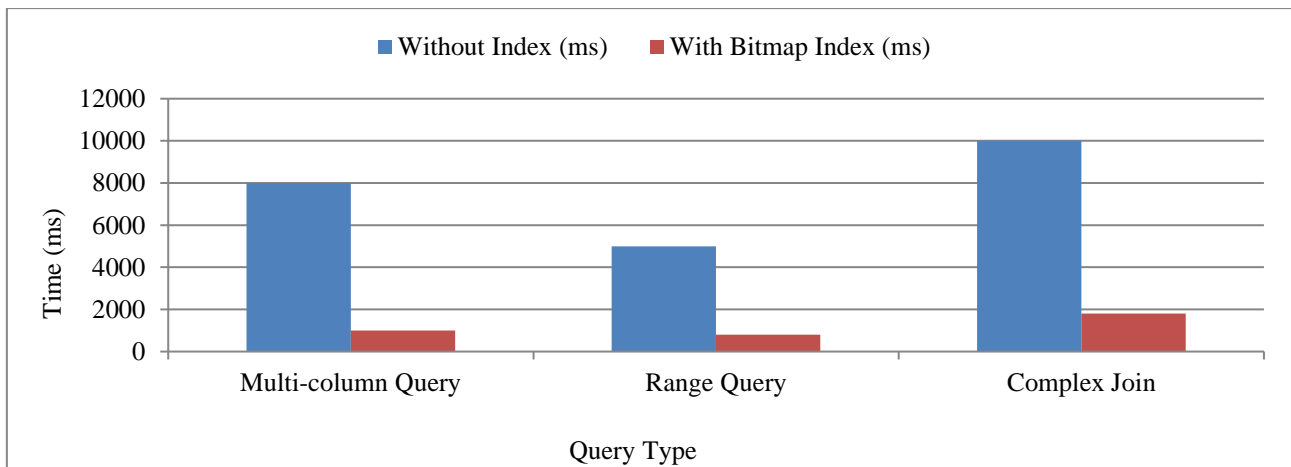


Fig. 3 Performance with and without Bitmap Index

5. Discussion

The findings highlight the importance of selecting the appropriate indexing technique based on the specific use case and query patterns.

5.1. B-tree Indexes

B-trees offer balanced performance across a variety of query types, making them a versatile choice for general-purpose indexing. Their logarithmic time complexity ensures scalability as datasets grow, and they are particularly effective in handling range and complex queries.

5.2. Hash Indexes

Hash indexes provide the fastest performance for equality searches but are limited by their inability to support range queries. They are best suited for use cases where queries are predominantly equality-based and where range queries are rare.

5.3. Bitmap Indexes

Bitmap indexes are highly effective for complex queries involving multiple columns and Boolean operations,

particularly in environments with low cardinality data. However, their higher storage requirements must be taken into account when deploying them in large-scale systems.

6. Conclusion

Indexing is an essential component of relational database optimization, offering significant improvements in query performance. This study provides practical insights into the benefits and limitations of various indexing techniques, guiding database administrators and developers in their efforts to design efficient and responsive databases.

Recommendations

- Analyze the types of queries most frequently executed to determine the most suitable indexing strategy.
- For queries involving multiple columns, consider using composite indexes to enhance performance.
- Continuously monitor query performance and adjust indexing strategies as data volumes and query patterns evolve.

References

- [1] Rudolf Bayer, and Edward McCreight, "Organization and Maintenance of Large Ordered Indexes," *Acta Informatica*, vol. 1, pp. 173-189, 1972. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [2] C. J. Date, *An Introduction to Database Systems*, 7th ed., Addison-Wesley, pp. 1-938, 2000. [[Google Scholar](#)] [[Publisher Link](#)]
- [3] Michael Stonebraker et al., "*The Design and Implementation of INGRES*," University of California, Berkeley, Technical Report, 1976. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [4] MySQL 8.0 Reference Manual: Including MySQL NDB Cluster 8.0, MySQL, 2024. [Online]. Available: <https://dev.mysql.com/doc/refman/8.0/en/>
- [5] Joseph M. Hellerstein, Michael Stonebraker, and James Hamilton, "Architecture of a Database System," *Foundations and Trends in Databases*, vol. 1, no. 2, pp. 141-259, 2007. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]