

QUERY OPTIMIZATION FOR BIG DATA ANALYTICS

Manoj Muniswamaiah, Tilak Agerwala and Charles Tappert

Seidenberg School of CSIS, Pace University, White Plains, New York

ABSTRACT

Organizations adopt different databases for big data which is huge in volume and have different data models. Querying big data is challenging yet crucial for any business. The data warehouses traditionally built with On-line Transaction Processing (OLTP) centric technologies must be modernized to scale to the ever-growing demand of data. With rapid change in requirements it is important to have near real time response from the big data gathered so that business decisions needed to address new challenges can be made in a timely manner. The main focus of our research is to improve the performance of query execution for big data.

KEYWORDS

Databases, Big data, Optimization, Analytical Query, Data Analysts and Data Scientists.

1. INTRODUCTION

Big data analytics is a process of gathering and analyzing data which is immense in volume, variety and velocity for making informed business decisions and take appropriate actions. There are different databases with varied data models to store and query big data: Columnar databases are used for read heavy analytical queries; online transactional processing databases are used for quicker writes and consistency; NoSQL data stores provide high read and write scalability. They are used in parallel processing of large volume of data and for horizontal distributed scaling; HTAP is a single database which can perform both OLTP and OLAP by utilizing in-memory computation. It is a hybrid of both OLTP and columnar database which can also be used for machine learning applications [1].

There are various data stores that are designed and developed to handle specific needs of big data and for optimizing performance. Relational databases can store and process structured data efficiently but its lacks scalability, elasticity needed to handle the inflow of big data. Its performance also decreases with read heavy queries. Similarly, columnar databases facilitate faster reads. This makes it a better choice for analytical query processing. NoSQL data stores provide horizontal scaling which makes it easier to add more resources on demand and thus are specialized to handle unstructured data. The processed data is then stored in different databases to be used by analysts for their business needs [1].

Performance optimization and different data models are important for data intensive applications. The main challenge is to build data pipelines that are scalable, interactive, efficient and fault tolerant. Data engineers optimize and maintain these data pipelines to ensure smooth data flow which are crucial for the performance of the applications. Data engineers primary job is to partition, normalize, index the base tables in order to improve the performance of the queries which are used for dashboards and reports. These optimized copies would be stored in different databases based on their data models for querying and quicker response. In this research we want to automate this process where if data scientists issues a query against any desired database the

query framework would be able to detect the optimized copies created and stored by data engineers and execute the query against optimized copy rather than the base table which would improve the performance of the query response.

2. BACKGROUND

There are several techniques to improve query performance and response time like partitioning the base table, indexing on required columns, materialization and creating OLAP cubes. Indexing helps in faster reads and retrieval of data. It is similar to data dictionary lookup, B-tree indexing keeps the data sorted and allows for sequential access of the data [2]. Consistency and freshness of the optimized copies is maintained by updating them whenever the base table changes.

BigDAWG, is a polystore framework implemented to support multiple databases compactible with different data models. The main feature of BigDAWG is to provide location independence which routes the queries to the desired databases and semantic completeness which lets the queries to make use of the native database features effectively. Shims translate incoming queries into their respective native database queries to take the advantages of the database features. One of the important features provided by BigDAWG is casting where data from one island is been converted into a model suitable for another island to enable inter-island queries and faster data migration between their underlying databases. The optimizer component parses the request, checks for syntax validation and creates a query plan to be executed by the engines. The monitor component uses statistics gathered from the previous queries to pick up the best query plan for execution. The migrator component is responsible for the movement of data between the data engines [3].

Apache Kylin is an open source distributed online analytic processing engine developed by eBay inc. It is built to support both ROLAP and MOLAP analytics. It provides SQL interface and sub-second query latency for large datasets which boosts its query performance. OLAP cubes are built offline using map reduce process. It reads data from source table and later uses map reduce to build cuboids of all combinations at each level. The cubes are generally stored in columnar database for faster reads. There are multiple cubing algorithms used for implementation. Recently Apache Spark is been used to speed up the cube build process which is stored in HBase data store. When users issue a query they are routed to be executed against the prebuilt cubes for quicker response, if the desired cube does not exists then the query would be executed against the Hadoop data [4].

Myria is an academia analytical engine been developed and provided as a cloud service. MyriaX is a share-nothing relational analytical query execution engine which efficiently executes the queries. It supports federated analytics and alleviates big data analysis. MyriaL is the query language supported by Myria for querying [5]

Apache Hive supports adhoc-queries and is used for batch processing of big data. It converts SQL-like queries in to map reduce jobs and executes the queries. HiveQL is the query language used for processing of the data [6].

Kodiak is an analytical distributed data platform which uses materialized views to process analytical queries. Various levels of materialized views are created and stored over petabytes of data. Kodiak platform can maintain these views efficiently and update them automatically. In Kodiak the query execution was three times faster and also used less resources [7].

Apache Lens is another open source framework which tries to provide a unified analytical layer

of Hadoop and databases ecosystem using a REST server and query language CubeQL to store and process the data cubes [8].

Analytical queries whose aggregates have been stored as OLAP cubes are used in reports and dashboards. These cubes often need to be updated with latest aggregates. Multiple databases are used within an organization for various tasks to store and retrieve the data. Data engineers implement data pipelines to optimize datasets which would be later used by data analysts and data scientists for research. Creating optimized copies involves partitioning and indexing of the base tables. These optimized copies would later be stored in different databases for querying.

Our research focus and solution is to implement a query framework that routes the data analysts and data scientists queries to the optimized copies created by data engineers which are stored, maintained, updated automatically to achieve better query performance and reduce response time.

3. QUERY OPTIMIZER

Apache calcite framework is an open source database querying framework which uses relational algebra for query processing. It parses the incoming coming query and converts them to logical plans and later various transformations would be applied to convert this logical plan into an optimized plan which has low cost in execution. This optimized logical plan would be converted into physical plan to be executed against the databases by using traits. Query optimizer eliminates logical plans which increase cost of the query execution based on cost model been defined. Apache Calcite Schema contains details about the data formats present in the model which is used by schema factory to create schema [9].

The query optimizer applies the planner rules to the relational node and generates different plans with reduced cost by retaining the original semantics of the query. When a rule matches a pattern query optimizer executes the transformations by substituting the subtree into the relation expression and also preserves the semantics of it. These transformations are specific to each database. The metadata component provides the query optimizer with details of the overall cost of the execution of the relational expression and also the degree of parallelism that can be achieved during execution [9].

The query optimizer uses cost-based dynamic programming algorithm which fires rules in order to reduce the cost of the relational expression. This process continues until the cost of the relational expression is not improved subsequently. The query optimizer takes into consideration CPU cycles, memory been used and IO resource utilization cost to execute the relational expression [9].

One of the techniques which is used to improve query processing is to use the optimized copies been created by data engineers. The query optimizer needs to have the ability to make use of these optimized copies to rewrite the incoming queries. Optimizer does this by substituting part of the relational expression tree with optimized copies which it uses to execute the query and return the response.

Algorithm

Optimization of relational expression R:

1. Register the relational expression R.
 - a. Check for the existence of appropriate optimized copy which can be used to substitute

- the relational expression R.
 - b. If the optimized copy exists, then register the new relational expression R1 of the optimized copy.
 - c. Trigger the transformation rules on relational expression R1 for cost optimization.
 - d. Obtain the best relational expression R1 based on cost and execute it against the database.
2. If the relational expression R cannot be substituted with an existing optimized copy
 - a. Trigger the transformation rules on relational expression R.
 - b. Obtain the best relational expression based on cost and execute it against the database[10].

4. ARCHITECTURE

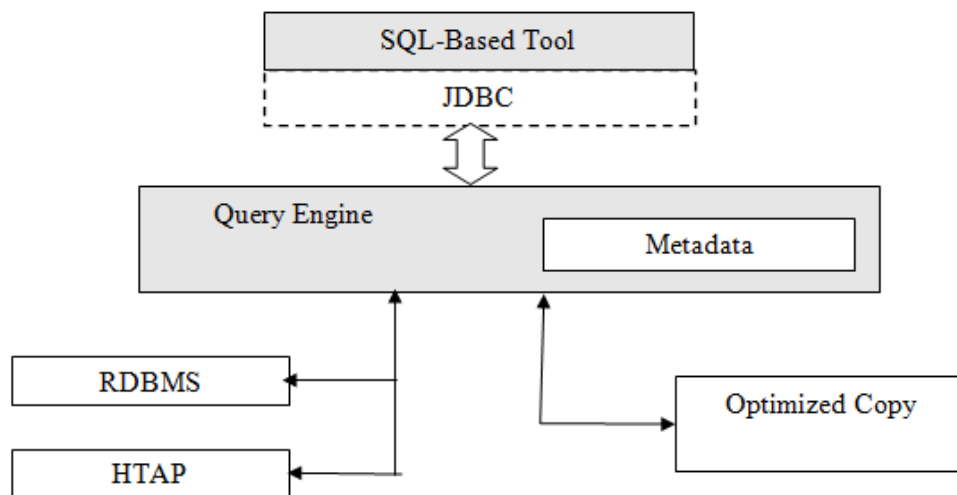


Figure 1. Architecture of the SQL Framework

The analytical queries been executed would be parsed and different logical plans would be generated. The query parser determines if the parsed relational expression can be substituted with the registered optimized copies been created by data engineers automatically. It executes various rules to obtain relational expression with minimal cost. The query would then be rewritten to be executed against the optimized copy and the results would be returned [10].

SQL-Tool: Includes any tool which can be integrated using JDBC connection and execute SQL analytical queries.

Query Engine: Apache Calcite open source framework that includes query optimizer been extended to include optimized copies.

Metadata: Contains information about the schema and the features of the native databases.

Optimized copy: Optimized tables created by the data engineers.

RDBMS: Includes any relational database to store structured data.

HTAP: Hybrid database which has the scalability of NoSQL data stores.

5. EVALUATION

Analytical queries help in data driven decision making process. In this paper we used NYC Taxi and Limousine Commission (TLC) datasets provided under the authorization of Taxicab & Livery Passenger Enhancement Programs consisting of the green and yellow taxi data [11].

The following experimental setup was made to benchmark and evaluate the query optimizer performance to find the optimized copy of the data and substitute it in the query plan during run time. Data engineers usually store these optimized copies in the columnar database for faster read access. The tables and data used for the query evaluation was obtained from NYC taxi dataset [11]. Data was stored in Mysql [12], Splice Machine [13] and Vertica database [14].

Evaluation was obtained based on the following setup

- a.) Base tables had rows ranging from ~10,000,000 to ~20,000,000
- b.) Optimized copy tables had rows ranging from ~1,000 to ~4,000,000
- c.) Mysql, Splice Machine and Vertica database was running on single node instance with Intel Quad Core Xeon 3.33GHz, 24 GB RAM and 1TB HDD.
- d.) Base table data was stored in HDFS [15].
- e.) Optimized copies were stored in Vertica database.
- f.) SQL query optimizer used cost based dynamic algorithm to substitute optimized copy in the query plan.

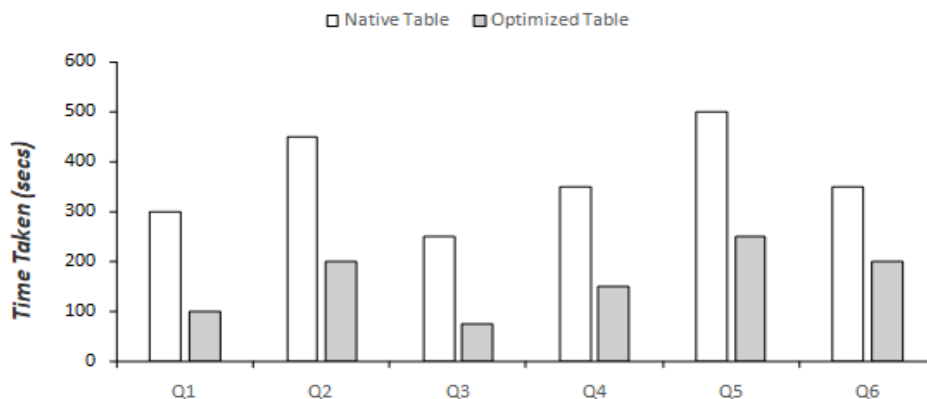


Figure 2. Query Response Time

The two bar graphs show the analytical queries been executed by the query optimizer against the base table and the optimized copy. The query optimizer was successfully able to substitute the optimized copy in the query plan when it existed and improved the performance of the query and its response time.

6. CONCLUSION

In this research we were able to make extensions to cost based query optimizer to make use of the optimized copies to obtain an improved query response time and performance. Various analytical queries were executed to measure the results obtained. The query optimizer was successfully able to substitute the optimized copies when existed during the runtime and modify the incoming query to execute against it. Part of the future work is to extend this query optimizer to the cloud databases.

REFERENCES

- [1] Duggan, J., Elmore, A. J., Stonebraker, M., Balazinska, M., Howe, B., Kepner, J., et al. (2015). The BigDAWG Polystore System. *ACM Sigmod Record*, 44(3)
- [2] V. Srinivasan and M. Carey. Performance of B-Tree Concurrency Control Algorithms. In *Proc.ACM SIGMOD Conf.*, pages 416–425, 1991
- [3] A. Elmore, J. Duggan, M. Stonebraker, M. Balazinska, U. Cetintemel, V. Gadepally, J. Heer, B. Howe, J. Kepner, T. Kraska et al., “A demonstration of the bigdawg polystore system,” *Proceedings of the VLDB Endowment*, vol. 8, no. 12, pp. 1908–1911, 2015
- [4] <http://kylin.apache.org>
- [5] D. Halperin et al. Demonstration of the myria big data management service. In *SIGMOD*, pages 881–884, 2014.
- [6] Fuad, A., Erwin, A. and Ipung, H.P., 2014, September. Processing performance on Apache Pig, Apache Hive and MySQL cluster. In *Information, Communication Technology and System (ICTS), 2014 International Conference on* (pp. 297-302). IEEE.
- [7] Liu, Shaosu, et al. "Kodiak: leveraging materialized views for very low-latency analytics over high-dimensional web-scale data." *Proceedings of the VLDB Endowment* 9.13 (2016): 1269-1280
- [8] <https://lens.apache.org/>
- [9] <https://calcite.apache.org/>
- [10] Muniswamaiah, Manoj & Agerwala, Tilak & Tappert, Charles. (2019). Query Performance Optimization in Databases for Big Data. 85-90. 10.5121/csit.2019.90908.
- [11] <https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page>
- [12] Luke Welling, Laura Thomson, PHP and MySQL Web Development, Sams, Indianapolis, IN, 2001
- [13] <https://www.splicemachine.com/>
- [14] C. Bear, A. Lamb, and N. Tran. The vertica database: Sql rdbms for managing big data. In *Proceedings of the 2012 workshop on Management of big data systems*, pages 37–38. ACM, 2012
- [15] Cong Jin, Shuang Ran, "The research for storage scheme based on Hadoop", *Computer and Communications (ICCC) 2015 IEEE International Conference on*, pp. 62-66, 2015.