

# Optimizing The Performance Of HADOOP Using Index-Based Join Operation

**Anshul Shrivastava**

*Dept. of Computer Science & Engineering  
Global Engineering College  
Jabalpur, Madhya Pradesh, India*

**Deepak Paranjpe**

*Asst. Professor, Dept. of Computer Science &  
Engineering  
Global Engineering College  
Jabalpur, Madhya Pradesh, India*

**Abstract** - Data refers to technologies and initiatives that involve data that is too diverse, fast-changing or massive for conventional technologies, skills and infrastructure to address efficiently is called Big Data. Hive is a data warehouse infrastructure tool, well suited for query processing and data analysis. Hive is gaining popularity for its SQL like query language HiveQL and for supporting majority of the SQL operations in relational database management systems (RDBMS). Being the expensive operation in RDBMS, join has been the focus of many query optimization techniques to improve performance of database systems. We investigate such techniques for join operations in Hive and develop a two-way join algorithm for queries in HiveQL. When a query requires only a small subset of data selected by a predicate in the WHERE clause, the brute-force method which scans the entire tables results in poor performance for redundant disk I/Os, and irrelevant maps initiation in case the query is issued using the MapReduce.

In this work, we implement the proposed index-based join technique and integrate it in Hive. To add our extension, we obtain Hive architecture details by reverse engineering the code and map our design to the conceptual optimization flow. To evaluate the performance, after setting up the environment, we run relevant test queries on datasets generated using the industry standard benchmark, TPC-H. Our results indicate significant performance gain over relatively large data or highly selective queries.

**KEYWORDS:** Big Data, Hive, Hadoop, Indexing, Join.

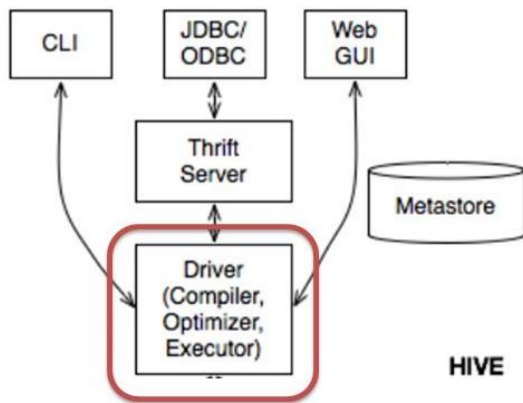
## I. INTRODUCTION

Data is growing at a huge speed making it difficult to handle such large amount of data (Giga-bytes). The main difficulty in handling such large amount of data is because that the volume is increasing rapidly in comparison to the computing resources. The Big data term which is being used now a days is kind of misnomer as it points out only the size of the data not putting too much of attention to its other existing properties [10]. The characteristics of Big Data can be broadly divided into three Vs i.e. Volume, Velocity and Variety and. Volume refers to the size of the data. While Velocity tells about the pace at which data is generated; Variety tells us about the complexity and structure of data and different ways of interpreting it [11].

Join is an expensive operation in databases, which depending on the predicate, data, etc., allows information from various relations to be “combined”. It also provides more data analysis and mining tasks important in the context of business intelligence for finding interesting and useful patterns in large amount of data. Therefore, improving various join operations can result in significant performance improvement. In relational databases, efficient join operations are supported through indexing without which the brute-force scan of the entire table is hopeless for large data. This is more needed when a small fraction of the tuple participate in a join operation. Two major factors that influence the performance of index based join operations in Hive includes very high data volume and low index maintenance cost [1]. Though Hive is expected to work well with vast amount of data, indexing can further improve the performance by reducing the amount of data accessed from the contributing tables. Having infrequent updates, as a characteristic of big data, makes the cost of index maintenance of less importance or affordable. Additionally, the index types proposed and developed in Hive take up a pretty small space. This paper presents a solution approach to perform join with Map-Reduce type operations, over large sets of data stored in a Hadoop-based cloud. We have experimentally evaluated the performance of the proposed approach, which uses a recent indexing feature of Hive to improve performance over previous- indexed queries. affordable. Additionally, the index types proposed and developed in Hive take up a pretty small space.

## II. HIVE

Hive is data warehouse software which is used for facilitates querying and managing large data sets residing in distributed storage. Hive language almost look like SQL language called HiveQL [4]. Hive is designed to enable easy data summarization. Hive also allows traditional map reduce programs to customize mappers and reducers when it is inconvenient or inefficient to execute the logic in HiveQL.



**Figure 1: Hive Component**

The first component is shell; Shell is the command line interface. It allows interactive queries like My SQL shell connected to database. Also supports web and JDBC clients [5]. Driver, compiler and execution engine take the HiveQL scripts and run in Hadoop environment.

The second component driver, which receives the queries. This component implements the notion of session handles and provides execute and fetch APIs modeled on JDBC/ODBC interfaces of the table and partition metadata looked up from the metastore. The third component in hive is Execution engine which executes the execution plan created by the compiler. The plan is a DAG of stages. The execution engine manages the dependencies between these different stages of the plan and executes these stages on the appropriate system components. The last component is meta store that stores all the structure information of the various table and partitions in the warehouse including column and column type information, the serializes and deserialize necessary to read and write data and the corresponding HDFS files where the data is stored.

### III. RELATED WORK

We will present a few optimization techniques related to indexes in Hive. HIVE-1644 [6] is the implementation of processing the WHERE clause with the index. The new query replaces the table with the index table and looks for the address of the desired values. The relevant part of the predicate is the part that can be processed by the indexes, that is a conjunction of the binary expressions. The re-written query is compiled and the produced root tasks are added to the original query root tasks. Then the original query is executed over the intermediate results produced from the re-written query. All column references in HIVE-1644 must refer to the same table (no joins or sub-queries).

```
SELECT col_list
FROM tab_name
WHERE predicate;
```

and re-writes it into:

```
INSERT INTO intermediate
SELECT _BUCKETNAME, _OFFSETS
FROM tab_index
WHERE relevent_part_of_the_predicate;
```

Antony, S., Chakka, 2010[1] proposed accelerates queries with GROUP BY clauses here are a number of conditions to be met in HIVE-1694: the FROM clause must have only one table (no joins) in the query; there should be only one COUNT (index\_key) function in the SELECT clause; and all column references must be in the index key.

In another research, Wang et Al 2010 [7] integrated indexing with a B+ tree structure into map-reduce framework. In this work, given a query, the index is accessed twice to locate the start point and the end point in the leaves. The nodes between these two positions satisfy the query. Map jobs are generated and attached to blocks of data covered between the start point and the end point. Each map first scans the index and then retrieves the records using the offset. Gruenhe id, work proposed storing column-level meta-data in Hive tables to benefit from during query execution [8]. Column-level statistics or more specifically, histograms that exhibit value distribution within a table provide more accurate information than just the table size to estimate the output size. A new table is added to Hive meta-store that holds the number of distinct values, number of null value, min and max values and most frequent values as its fields. In presence of column statistics, an index-based join can determine whether it is an optimal approach before execution.

### IV. Index-based Join approach

Single tables are involved in existing indices in Hive. A join index is a pre-computed access structure that maintains pairs of identifiers of tuples from two or more relations that would match in case of a join in RDBMS. This approach would be a suitable optimization approach in Hive where tables are updated infrequently. Index join concept is based on the fact that is to keep unique identifiers of the matched tuples in the same structure and cluster them on either of the unique identifiers of both tables. The current implementation of Hive does not support the concept of primary keys [9] which are considered the unique identifiers of tuples in RDBMSs.

The aim of my work is to accelerate a two-way join query created in HiveQL as shown below:

```
SELECT col_list
FROM table1 JOIN table2
ON (table1.col1 = table2.col1)
[WHERE ...]
[GROUP BY...];
```

In the above mentioned query “WHERE” and “GROUP BY” clauses are optional. The same queries can be applied for joining “n” number of tables.

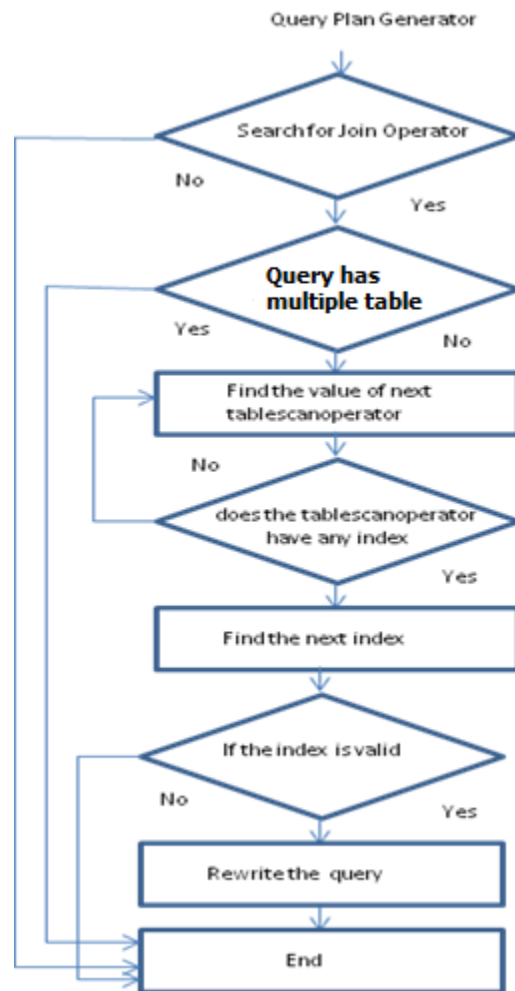
**A. Design**

When we execute simple query in Hive, it reads the whole dataset even if we have use ‘where ‘clause filter. This becomes a bottleneck for running Map-Reduce jobs over a huge table. We can overcome this problem by using partitioning in Hive. By using automatic partition method when the table is created. In Hive’s implementation of partitioning, data within a table is split across various small partitions. When the query is executed, only the required partitions of the table are execute, thereby minimizing the I/O and time required by the query. Because when external table is declared, default table path is changed to specified location in hive metadata which contains in meta store, but about partition, nothing is changed, so, we must manually add those metadata.

The proposed research work can be demonstrated by the following:

A Search for a JoinOperator is done by optimizer. If this step is omitted we can perform optimization for any query. Now query is examined by optimizer for a two-way join. Further we get operator TableScanOperator that points to the table that has to be manipulated and verify that the table contain an index and check for its validity. The index is valid if it is compact index and it includes all the partitions of the table .If all the condition are fulfilled then the optimizer re-write the query:

```
SELECT col_list
FROM index_table JOIN table2
ON (tab1.col1 = tab2.col1)
[WHERE]
[GROUP BY];
```



**Figure: 2 Query Plan Generator Flow of Index-based approach**

Otherwise this flow ends which means query is not executed successfully.

Any of the table (whichever that has the index) is replaced by its corresponding index table. This means that table must be removed from every internal data structure in the operator DAG and the new table must be added instead.

Hive query optimization is conforms by given flow show in figure 2.

**B. Results**

We used the standard benchmark TPC-H version 2.14.4 to generate data used in our experiments [10]. We considered only the supplier and nation tables.

**C. Test Queries**

- **SELECT** s\_acctbal,s\_s\_name,s\_s\_suppkey from supplier s **JOIN** nation n **ON**(s.s\_nationkey=n.n\_nationkey) .

- **SELECT** s.s\_acctbal,s.s\_name,s.s\_suppkey from supplier s **JOIN** nation n **ON** (s.s\_nationkey=n.n\_nationkey) group by s.s\_acctbal,s.s\_name,s.s\_suppkey **WHERE** c.c\_acctbal>1000.

- **SELECT** s.s\_acctbal,s.s\_name,s.s\_suppkey from supplier s **JOIN** nation **ON** (s.s\_nationkey=n.n\_nationkey) group by s.s\_acctbal,s.s\_name,s.s\_suppkey **GROUPBY** s.s\_acctbal,s.s\_name,s.s\_suppkey .

Data Size	Without index approach response time(s)			With index approach response time(s)		
	1 GB	3 GB	5 GB	1 GB	3 GB	5 GB
	50.60	147.66	265.00	32.11	101.56	210.54
	49.81	146.87	265.78	32.45	101.43	211.34
	49.43	150.98	265.65	32.12	101.24	210.11
	49.76	147.01	265.78	32.43	101.23	210.45
	49.76	147.00	265.71	32.10	101.45	210.23
<b>Avg.</b>	<b>49.87</b>	<b>147.90</b>	<b>265.58</b>	<b>32.24</b>	<b>101.38</b>	<b>210.53</b>

Table 3: Query1 Response time without index based approach /with index-based approach

Data Size	Without index approach response time(s)			With index approach response time(s)		
	1 GB	3 GB	5 GB	1 GB	3 GB	5 GB
	51.22	146.11	200.54	34.11	100.12	143.23
	52.56	142.25	201.56	32.34	100.34	143.22
	55.78	142.34	203.54	31.67	101.23	145.32
	52.44	141.26	203.55	30.00	101.45	142.23
	52.11	140.99	201.45	30.67	101.34	142.23
<b>Avg.</b>	<b>52.82</b>	<b>142.59</b>	<b>202.128</b>	<b>31.758</b>	<b>100.896</b>	<b>143.246</b>

Table 1: Query1 Response time without index based approach /with index-based approach

Data Size	Without index approach response time(s)			With index approach response time(s)		
	1 GB	3 GB	5 GB	1 GB	3 GB	5 GB
	54.22	141.11	193.54	32.11	110.12	133.23
	53.56	142.25	191.56	32.34	110.34	133.22
	53.78	142.34	191.54	31.67	111.23	135.32
	52.44	141.26	191.55	31.02	111.45	132.23
	52.11	140.99	191.45	31.23	111.34	132.23
<b>Avg.</b>	<b>53.11</b>	<b>141.59</b>	<b>191.928</b>	<b>31.674</b>	<b>110.896</b>	<b>133.246</b>

Table 2: Query1 Response time without index based approach /with index-based approach




Figure 3: Query1 response time without index based /with index-based approach on single node setup




Figure 4: Query2 response time without index based /with index-based approach on single node setup

**REFERENCES**

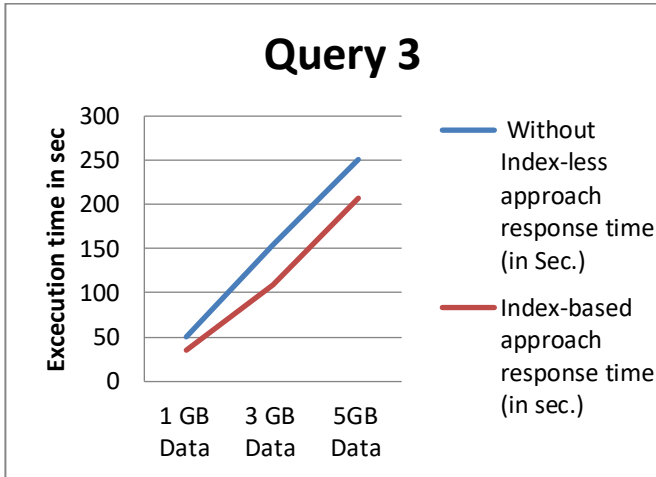


Table 3: Query1 Response time without index based approach /with index-based approach

**V. CONCLUSION AND FUTURE WORK**

The expensive operation in RDBMS, join has been the focus of many query optimization techniques to improve performance of database systems. In partitioning, we create a partition for each unique value of the column. We investigate such techniques for join operations in Hive and develop an index-based join algorithm for queries in HiveQL. Indexes have been around for long time and the benefit of using them is obvious though index size depends on the data distribution and the number of attributes for indexing, our experiments showed the Hive index space utilization is reasonable. With respect to accessing the index, current Hive indexes do not provide an instant access to values, which undoubtedly comes with heavy space overhead. What they offer instead is, scanning a huge amount of data is replaced with scanning a drastically small set of it that holds the desired values. Hive index maintenance cost is noticeably low, considering the infrequent updates and batch-mode data insertion as the characteristics of big data. The indexing technique in Hive is rather new and the progress has been limited to current index structure and also the query life cycle. As future work, first we plan to work on hash based indexing using bucket level because bucket is smallest data model in hive .

[1] Antony, S., Chakka, P., Jain, N., J., Liu, Murthy, R., Sarma, J. S., Thusoo, A., Zhang, N “Hive – A Petabyte Scale Data Warehouse Using Hadoop,” IEEE 26th Intl. Conf. Data Engineering (ICDE), Long Beach, CA, 2010, pp. 996 – 1005.

[2] Apache Hadoop [Online]. Available: <http://hadoop.apache.org/>

[3] Dean, J., Ghemawat, S. “MapReduce: Simplified Data Processing on Large Clusters,” Mag. Commun. ACM 50th anniversary, vol. 51, issue 1, 2008, pp.107-113

[4] <http://www.hadooppoint.com/introduction-hive/>

[5] Yue Liu1,6,7 , Songlin Hu1 “DGFIndex for Smart Grid: Enhancing Hive with a Cost-Effective multidimensional Range Index” 40th International Conference on Very Large Data Bases, September 1st - 5th 2014, Hangzhou, China.

[6] ANTLR [Online]. Available: <http://www.antlr.org/>

[7] An, M., Wang, W., Wang, Y., “Using Index in the MapReduce Framework, ”, 12th Intl. Asia Pacific Web Conf. (APWEB),Beijing, China, 2010, pp. 52-58

[8] Dean, J., Ghemawat, S. “MapReduce: Simplified DataProcessing on Large Clusters,” Mag. Commun. ACM 50th anniversary, vol. 51, issue 1, 2008, pp.107-113

[9]Capriolo, E., Rutherglen, J., Wampler, D. Programming Hive: Data Warehouse and Query Language for Hadoop, 1st ed, O'Reilly Media, 2012

[10] TPC-H[Online]. <http://www.tpc.org/tpch/>

[11]HIVE 1694[Online]. Available: <https://issues.apache.org/jira/browse/HIVE-1694>

[12] Hive index design doc [Online]. Available: <https://cwiki.apache.org/confluence/display/Hive/IndexDe>

[13] Hive JIRA [Online]. Available: <https://issues.apache.org/jira/browse/HIVE>

[14] HIVE-1644 [Online]. Available: <https://issues.apache.org/jira/browse/HIVE-1644>

[15] N. Jain, L. Tang, “Join strategies in Hive”, Facebook, Rep. Hadoop summit 2011, 2011 [Online].

[16] Li, Z., Ross, K. A. “Fast joins using join indices”, in The International Journal on Very Large Data Bases, vol. 8, issue 1, 1999, pp.1–24

[17] Lou, W., Ren, K., Wang, C., Wang, Q. Privacy-Preserving Public Auditing for Storage Security in Cloud Computing, Proc. 30th IEEE Int'l Conf. Computer Communications (INFOCOM 10), IEEE Press, San Diego, CA, 2010, pp. 525–533.

[18] S. Madden: “From Databases to Big Data,” IEEE Internet Computer., vol.16, issue 3, pp. 4-6, May-June, 2012

[19] MapReduce Tutorial [Online]. Available (date): [http://hadoop.apache.org/docs/mapreduce/r0.22.0/mapred\\_tutorial.html](http://hadoop.apache.org/docs/mapreduce/r0.22.0/mapred_tutorial.html)

[20]MongoDB[Online]. Available: <http://www.mongodb.org/>

[21]Neo4j[Online]. Available(write date): <http://www.neo4j.org/>

[22] Gruenheid, A., Mark, L., Omnecinski, E. “Query Optimization using column statistics in Hive,” in Proc. 15th Symp. Intl. Database Engineering & Applications (IDEAS), Lisbon, Portugal, 2011, pp. 97-105, 2011



## **International Journal of Recent Development in Engineering and Technology**

**Website: [www.ijrdet.com](http://www.ijrdet.com) (ISSN 2347-6435(Online) Volume 7, Issue 4, April 2018)**

- [23] Dean, J., Ghemawat, S. "MapReduce: Simplified Data Processing on Large Clusters," Mag. Commun. ACM 50th anniversary, vol. 51, issue 1, 2008, pp.107-113
- [24] Eaton, C., Deroos, D., Deutsch, T., Lapis, G., Zikopoulos, P. Understanding Big data: Analytics for Enterprise Class Hadoop and Streaming Data, 1st ed, McGraw, 2011
- [25] Garcia-Molina, H., Ullman, J., Widom, J. Database Systems: The Complete Book, 1st ed, Upper Saddle River, NJ, Prentice Hall Inc., 2002
- [26] Gilbert, S., Lynch, N. A. "Brewer's Conjecture and the Feasibility of Consistent, Available, Partition-Tolerant Web Services", in Newslett. ACM SIGACT, vol. 33, issue 2, pp. 51-59, June 2002
- [27] Grance, T., Mell, P. "The NIST Definition of Cloud Computing," NIST. Gaithersburg, MD, Rep. Recommendations of the National Institute of Standards and 107 Technology, 2011
- [28] Yin Huai, Ashutosh Chouhan "Major Technical Advancements in Apache Hive" SIGMOD 2014, June .
- [29] H. V. Jagadish : "Big Data & Science Myths and Reality" Science Direct June 2015.