

Comparative Study Parallel Join Algorithms for MapReduce environment

A. Pigul
m05pay@math.spbu.ru
Saint Petersburg State University

Abstract. There are the following techniques that are used to analyze massive amounts of data: MapReduce paradigm, parallel DBMSs, column-wise store, and various combinations of these approaches. We focus in a MapReduce environment. Unfortunately, join algorithms is not directly supported in MapReduce. The aim of this work is to generalize and compare existing equi-join algorithms with some optimization techniques.

Key Words: parallel join algorithms, MapReduce, optimization.

1. Introduction

Data-intensive applications include large-scale data warehouse systems, cloud computing, data-intensive analysis. These applications have their own specific computational workload. For example, analytic systems produce relatively rare updates but heavy select operation with millions of records to be processed, often with aggregations.

Applications for large-scale data analysis use such techniques as parallel DBMS, MapReduce (MR) paradigm, and columnar storage. Applications of this type process multiple data sets. This implies need to perform several join operation. It's known join operation is one of the most expensive operations in terms both I/O and CPU costs.

Unfortunately, join algorithms is not directly supported in MapReduce. There are some approaches to solve this problem by using a high-level language PigLatin, HiveQL for SQL queries or implementing algorithms from research papers. The aim of this work is to generalize and compare existing equi-join algorithms with some optimization techniques.

This paper is organized as follows the section 2 describe state of the art. Join algorithms and some optimization techniques were introduced in 3 section. Performance evaluation will be described in 4 section. Finally, future direction and some discussion of experiments will be given.

2. Related Work

2.1. Architectural Approaches

Column storage is one of the architectural approaches to store data in columns, that the values of one field are stored physically together in a compact storage area. Column storage strategy improves performance by reducing the amount of unnecessary data from disk by excluding the columns that are not needed. Additional gains may be obtained using data compression. Storage method in columns outperforms row-based storage for workloads typical for analytical applications, which are characterized by heavy selection operation from millions of records, often with aggregation and by infrequent update operation. For this class of workloads I/O is major factor limited the performance. Comparison of column-wise and row-wise stores approaches is presented in [1].

Another architectural approach is a software framework MapReduce. Paradigm MapReduce was introduced in [11] to process massive amounts of unstructured data.

Originally, this approach was contrasted with a parallel DBMS. Deep analysis of the advantages and disadvantages of these two architectures was presented in [25,10].

Later, hybrid systems appeared in [9, 2]. There are three ways to combine approaches MapReduce and parallel DBMS.

- MapReduce inside a parallel DBMS. The main intention is to move computation closer to data. This architecture can be exemplified with hybrid database Greenplum with MAD approach [9].
- DBMS inside MapReduce. The basic idea is to connect multiple single node database systems using MapReduce as the task coordinator and network communication layer. An example is a hybrid database HadoopDB [2].
- MapReduce aside of the parallel DBMS. MapReduce is used to implement an ETL produced data to be stored in parallel DBMS. This approach is discussed in [28] Vertica, which also supports the column-wise store.

Another group of hybrid systems combines MapReduce with column-wise store. MapReduce and column-wise store are effective in data-intensive applications. Hybrid systems based on this two techniques may be found in [20,13].

2.2. Algorithms for Join Operation

Detailed comparison of relational join algorithms was presented in [26]. In our paper, the consideration is restricted to a comparison of joins in the context of MapReduce paradigm.

Papers which discuss equi-join algorithms can be divided into two categories which describe join algorithms and multi join execution plans.

The former category deals with design and analyses join algorithm of two data sets. A comparative analysis of two-way join techniques is presented in [6, 4, 21]. The cost model for two-way join algorithms in terms of cost I/O is presented in [7, 17].

The basic idea of multi-way join is to find strategies to combine the natural join of several relations. Different join algorithms from relation algebra are presented in [30]. The authors introduce the extension of MapReduce to facilitate implement relation operations. Several optimizations for multi-way join are described in [3, 18]. Authors introduced a one-to-many shuffling strategy. Multi-way join optimization for column-wise store is considered in [20, 32].

Theta-Joins and set-similarity joins using MapReduce are addressed in [23] and [27] respectively.

2.3. Optimization techniques and cost models

In contrast to the sql queries in parallel database, the MapReduce program contains user-defined map and reduce functions. Map and reduce functions can be considered as a black-box, when nothing is known about these functions, or they can be written on sql-like languages, such as HiveQL, PigLatin, MRQL, or sql operations can be extracted from functions on semantic basis. Automatic finding good configuration settings for arbitrary program offered in [16]. Theoretical designing cost models for arbitrary MR program for each phase separately presented in [15]. If the MR program is similar to the semantics of SQL, it allows us to construct a more accurate cost model or adapt some of the optimization techniques from relational databases. HadoopToSQL [22] allows to take advantage of two different data storages such as SQL database and the text format in MapReduce storage and to use index at right time by transforming the MR program to SQL. Minimal system [17] uses static analysis for detection and exploiting selection, projection and data compression in MR programs and if needed to employ B+ tree index.

New SQL-like query language and algebra is presented in [12]. But they are needed cost model based on statistic. Detailed construction of the model to estimate the I/O cost for each phase separately is given in [24]. Simple theoretical considerations for selecting a particular join algorithm are presented in [21]. Another approach [7] for selecting join algorithm is to measure the correlation between the input size and the join algorithm execution time with fixed cluster configuration settings.

3. Join algorithms and optimization techniques

In this section we consider various techniques of two-way joins in MapReduce framework. Join algorithms can be divided into two groups: Reduce-side join and Map-side join. The pseudo code presented in Listings, where R – right dataset, L – left dataset, V – line from file, Key – join key, that was parsed from a tuple, in this context tuple is V.

3.1. Reduce-Side join

Reduce-side join is an algorithm which performs data pre-processing in Map phase, and direct join is done during the Reduce phase. Join of this type is the most general without any restriction on the data. Reduce-side join is the most time-consuming, because it contains an additional phase and transmits data over the network from one phase to another. In addition, the algorithm has to pass information about source of data through the network. The main objective of the improvement is to reduce the data transmission over the network from the Map task to the Reduce task by filtering the original data through semi-joins. Another disadvantage of this class of algorithms is the sensitivity to the data skew, which can be addressed by replacing the default hash partitioner with a range partitioner.

There are three algorithms in this group:

- General reducer-side join,
- Optimized reducer-side join,
- the Hybrid Hadoop join.

General reducer-side join is the simplest one. The same algorithms are called Standard Repartition Join in [6]. The abbreviation is GRSJ and pseudo code is presented in Listing 1.

This algorithm has both Map and Reduce phases. In the Map phase, data are read from two sources and tags are attached to the value to identify the source of a key/value pair. As the key is not effecting by this tagging, so we can use the standard hash partitioner. In Reduce phase, data with the same key and different tags are joined with nested-loop algorithm. The problems of this approach are that the reducer should have sufficient memory for all records with a same key; and the algorithm sensitivity to the data skew.

```

Map (K: null, V from R or L)
  Tag = bit from name of R or L;
  emit (Key, pair(V,Tag));

Reduce (K': join key, LV: list of V with key K')
  create buffers Br and Bl for R and L;
  for t in LV do
    add t.v to Br or Bl by t.Tag;
  for r in Br do
    for l in Bl do
      emit (null, tuple(r.V,l.V));
  
```

Listing 1: GRSJ.

```

Map (K:null, V from R or L)
  Tag = bit from name of R or L;
  emit (pair(Key,Tag), pair(V,Tag));

Partitioner(K:key, V:value, P:the number of reducers)
  return hash_f(K.Key) mod P;

Reduce (K': join key, LV: list of V' with key K')
  create buffers Br for R;
  for t in LV with t.Tag corresponds to R do
    add t.v to Br;
  for l in LV with l.Tag corresponds to L do
    for r in Br do
      emit (null, tuple(r.V,l.V));

```

Listing 2: ORSJ.

Optimized reducer-side join enhances previous algorithm by overriding sorting and grouping by the key, as well as tagging data source. Also known as Improved Repartition Join in [6], Default join in [14]. The abbreviation is ORSJ. In Listing 2 pseudo code is shown. In the algorithm all the values of the first tag are followed by the values of the second one. In contrast with the General reducer-side join, the tag is attached to both a key and a value. Due to the fact that the tag is attached to a key, the partitioner must be overridden in order to split the nodes by the key only. This case requires buffering for only one of input sets.

Optimized reducer-side join inherits major disadvantages of General reducer-side join namely the transferring through the network additional information about the source and the algorithm sensitivity to the data skew.

The Hybrid join [4] combines the Map-side and Reduce-side joins. The abbreviation is HYB and Listing 3 describe pseudo code.

```

Job 1: partition the smaller file S
Map (K:null, V from S)
  emit (Key,V);
Reduce (K':join key, LV: list of V' with key K')
  for t in LV do
    emit (null, t);

Job 2: join two datasets
Map (K:null, V from B)
  emit (Key,V);
  init() //for Reduce phase
  read needed partition of output om Job 1;
  add it to hashMap(Key, list(V)) H;
Reduce (K':join key, LV: list of
  V' with key K')
  if(K' in H) then
    for r in LV do
      for l in H.get(K') do
        emit (null, tuple(r,l));

```

Listing 3: HYB

In Map phase, we process only one set and the second set is partitioned in advance. The pre-partitioned set is pulled out of blocks from a distributed system in the Reduce phase, where it is joined with another data set that came from the Map phase. The similarity with the Map-side join is the restriction that one of the sets has to be split in advance with the same partitioner, which will split the second set. Unlike Map-side join, it is necessary to split in advance only one set. The similarity with the Reduce-side join is that algorithm requires two phases, one of them for pre-processing of data and one for direct join. In contrast with the Reduce-side join we do not need additional information about the source of data, as they come to the Reducer at a time.

3.2. Map-Side join

Map-side join is an algorithm without Reduce phase. This kind of join can be divided into two groups. First of them is partition join, when data previously partitioned into the same number of parts with the same partitioner. The relevant parts will be joined during the Map phase. This map-side join is sensitive to the data skew. The second is in memory join, when the smaller dataset send whole to all mappers and bigger dataset is partitioned over the mappers. The problem with this type of join occurs when the smaller of the sets cannot fit in memory.

There are three methods to avoid this problem:

- JDBM-based map join,
- Multi-phase map join,
- Reversed map join.

Map-side partition join algorithm assumes that the two sets of data pre-partitioned into the same number of splits by the same partitioner. Also known as default map join. The abbreviation is MSPJ and Listing 4 describe pseudo code. At the Map phase one of the sets is read and loaded into the hash table, then two sets are joined by the hash table. This algorithm buffers all records with the same keys in memory, as is the case with skew data may fail due to lack of enough memory.

```

Job 1: partition dataset S as in HYB
Job 2: partition dataset B as in HYB
Job 3: join two datasets
  init() //for Map phase
  read needed partition of output file from Job 1;
  add it to hashMap(Key, list(V)) H;
Map(K:null, V from B)
  if (K in H) then
    for r in LV do
      for l in H.get(K) do
        emit(null, tuple(r,l));

```

Listing 4: MSPJ

```

Job 1: partition S dataset as in HYB
Job 2: partition B dataset as in HYB
Job 3: join two datasets
init() //for Map phase
    find needed partition SP of output file from Job 1;
    read first lines with the same key K2 from SP and add
    to buffer Bu;
Map(K:null, V from B)
    while (K > K2) do
        read T from SP with key K2;
        while (K == K2) do
            add T to Bu;
            read T from SP with key K2;
        if (K == K2) then
            for r in Bu do
                emit(null, tuple(r,V));

```

Listing 5: MSPMJ.

Map-side partition merge join is an improvement of the previous version of the join. The abbreviation is MSPMJ and pseudo code is presented in Listing 5. If data sets in addition to their partition are sorted by the same ordering, we apply merge join. The advantage of this approach is that the reading of the second set is on-demand, but not completely, thus memory overflow can be avoided. As in the previous cases, for optimization can be used the semi-join filtering and range partitioner.

In-Memory Join does not require to distribute original data in advance unlike the versions of map joins discussed above. The same algorithms are called Map-side replication join in [7], Broadcast Join in [6], Memory-backed joins [4], Fragment-Replicate join in [14]. The abbreviation is IMMJ. Nevertheless, this algorithm has a strong restriction on the size of one of the sets: it must fit completely in memory. The advantage of this approach is its resistance to the data skew because it sequentially reads the same number of tuples at each node. There are two options for transferring the smaller of the sets:

- using a distributed cache,
- reading from a distributed file system.

```

init() // for Map phase
    read S from HDFS;
    add it to hashMap(Key, list(V)) H;
map (K:null, V from B)
    if (K in H) then
        for l in H.get(K) do
            emit (null, tuple(v,l));

```

Listing 6: IMMJ

```

init() //for Map phase
    read S from HDFS;
    add it to hashMap(Key, list(V)) H;
map (K:null, V from S)
    add to hashMap(Key, V) H;
close() //for Map phase
    find B in HDFS
    while (not end B) do
        read line T;
        K = join key from tuple T;
        if (K in H) then
            for l in H.get(K) do
                emit(null, tuple(T,l));

```

Listing 7: REV.

The next three algorithms optimize the In-Memory Join for a case, when two sets are large and no of them fits into the memory.

JDBM-based map join is presented in [21]. In this case, JDBM library automatically swaps hash table from memory to disk.

```

The same as IMMJ, but H is implemented by HTree
instead of hashMap .

```

Listing 8: JDBM

```

For part P from S that fit into memory do IMMJ(P,B).

```

Listing 9: Multi-phase map join.

Multi-phase map join [21] is algorithm where the smaller of the sets is partitioned into parts that fit into memory, and for each part runs In-Memory join. The problem

with this approach is that it has a poor performance. If the size of the set, which to be put in the memory is increased twice, the execution time of this join is also doubled. It is important to note that the set, which will not be loaded into memory, will be read many times from the disk.

Idea of Reversed map join [21] approach is that the bigger of the sets, which is partitions during the Map phase, loading in the hash table. Also known as Broadcast Join in [6]. The abbreviation is REV. The second dataset is read from a file line by line and joined using a hash table.

3.3. Semi-Join

Sometimes a large portion of the data set does not take part in the join. Deleting of tuples that will not be used in join significantly reduces the amount of data transferred over the network and the size of the dataset for the join. This preprocessing can be carried out using semi-joins by selection or by a bitwise filter. However, these filtering techniques introduce some cost (an additional MR job), so the semi-join can improve the performance of the system only if the join key has low selectivity. There are three ways to implement the semi-join operation:

- a semi-join using bloom-filter,
- semi-join using selection,
- an adaptive semi-join.

Bloom-filter is a bit array that defines a membership of element in the set. False positive answers are possible, but there are no false-negative responses in the solution of the containment problem. The accuracy of the containment problem solution depends on the size of the bitmap and on the number of elements in the set. These parameters are set by the user. It is known that for a bitmap of fixed size m and for the data set of n tuples, the optimal number of hash functions is $k=0.6931*m/n$. In the context of MapReduce, the semi-join is performed in two jobs. The first job consists of the Map phase, in which keys from one set are selected and added to the Bloom-filter. The Reduce phase combines several Bloom-filters from first phase into one. The second job consists only of the Map phase, which filters the second data set with a Bloom-filter constructed in previous job. The accuracy of this approach can be improved by increasing the size of the bitmap. However in this case, a larger bitmap consumes more amounts of memory. The advantage of this method is its the compactness. The performance of the semi-join using Bloom-filter highly depends on the balance between the Bloom-filter size, which increases the time needed for its reconstruction of the filter in the second job, and the number of false positive responses in the containment solution. The large size of the data set can seriously degrade the performance of the join.

```

Job 1: construct Bloom filter
Map (K:null, V from L)
    Add Key to BloomFilter B1
close() //for Map phase
    emit(null, B1);

Reduce (K': key, LV) //only 1 Reducer
for l in LV do
    union filters by operation Or
close() // for Reduce phase
    write resulting filter into file;

Job 2: filter dataset
init() //for Map phase
    read filter from file in B1
Map (K:null, V from R)
    if (Key in B1) then
        emit (null, V);

Job 3: do join with L dataset and filtered dataset from
Job 2.

```

Listing 10: Semi-join using Bloom-filter

```

Job 1: find unique keys
Map (K:null, V from L)
    Create HashMap H;
    if (not Key in H) then
        add Key to H;
        emit (Key, null);

Reduce (K': key, LV) //only one Reducer
    emit (null,key);

Job 2: filter dataset
init() //for Map phase
    add to HashMap H unique keys from job 1;
Map (K:null, V from R)
    if (Key in H) then
        emit (null,V);

Job 3: do join with L dataset and filtered dataset from
Job 2.

```

Listing 11: Semi-join with selection.

Semi-join with selection extracts unique keys and constructs a hash table. The second set is filtered by the hash table constructed in the previous step. In the

context of MapReduce, the semi-join is performed in two jobs. Unique keys are selected during the Map phase of the first job and then they are combined into one file during the Map phase. The second job consists of only the Map phase, which filters out the second set. The semi-join using selection has some limitations. Hash table in memory, based on records of unique keys, can be very large, and depends on the key size and the number of different keys.

The Adaptive semijoin is performed in one job, but filters the original data on the flight during the join. Similar to the Reduce-side join at the Map phase the keys from two data sets are read and values are set equal to tags which identify the source of the keys. At the Reduce phase keys with different tags are selected. The disadvantage of this approach is that additional information about the source of data is transmitted over the network.

<p>Job 1: find keys which are present in two datasets</p> <p>Map (K:null, V from R or L) Tag = bit from name of R or L; emit (Key,Tag);</p> <p>Reduce (K': join key, LV: list of V with key K') Val = first value from LV; for t in LV do if (not Val==Val2) then emit (null, K');</p>	<p>Job 2: before joining it is necessary to filter the smaller dataset dataset by keys from the Job 1 that will be loaded into hash map. Then the bigger dataset is joined with filtered one</p>
--	--

Listing 12: Adaptive semi-join.

3.4. Range Partitioners

All algorithms, except the In-Memory join and their optimizations are sensitive to the data skew. This section describes two techniques of the default hash partitioner replacement.

A Simple Range-based Partitioner [4] (this kind similar to the *Skew join* in [14]) applies a range vector of dimension n constructed from the join keys before starting a MR job. By this vector join keys will be splitted into n parts, where n is the number of Reduce jobs. Ideally partitioner vector is constructed from the whole original set of keys, in practice a certain number of keys is chosen randomly from the data set. It is known that the optimal number of keys for the vector construction is equal to the square root of the total number of tuples. With a heavy data skew into a single key value, some elements of the vector may be identical. If the key belongs to multiple nodes, a node is selected randomly in the case of data on which to build

a hash table, otherwise the key is sent to all nodes (to save memory as a hash table is contained in the memory).

Virtual Processor Partitioner [4] is an improvement of the previous algorithm based on increasing the number of partition. The number of parts is specified multiple of the tasks number. The approach tends to load the nodes with the same keys uniformly (compared with the previous version). The same keys are scattered on more nodes than in the previous case.

```

//before the MR job starts
// optimal max = sqrt(|R|+|L|)
getSamples (Red:the number of reducers,
max: the max
           number of samples)
  C = max/Splits.length;
  Create buffer B;
  for s in Splits of R and L do
    get C keys from s;
    add it to B
  sort B;
  //in case simple range partitioner P == 1
  //in case virtual range partitioner P > 1
  for j<(Red*P) do
    T = B.length/(Red*P)*(j+1);
    write into file B[T];

Map(K:null, V from L or R)
Tag = bit from name of R or L;
read file with samples and add samples
to Buffer B;
//in case virtual partition it is needed to
// each index mod |Reducers|
Ind = {i: B[i-1] < Key <= B[i]}
// Ind may be array of indexes in skew
case
if (Ind.length >1) then
  if (V in L) then
    node = random(Ind);
emit (pair(Key, node), pair(V, Tag));
else
for i in Ind do
  emit (pair(Key, i), pair(V, Tag));
Partitioner (K:key, V:value, P:
the number of reducers)
return K.Ind;
Reducer (K': join key, LV: list of
V' with key K')
The same as GRSJ

```

Listing 13: The range partitioners.

3.5. Distributed cache

The advantage of using distributed cache is that data set are copied only once at the node. It is especially effective if several tasks at one node need the same file. In contrast the access to the global file system needs more communication between the nodes. Better performance of the joins without the cache can be achieved by increasing number of the files replication, so there's a good chance to access the file version locally.

3.6. Comparative analysis of algorithms

The features of join algorithms are presented in the Table 1. The approaches with pre-processing is good when data is prepared in advance for example come from another MapReduce job. Algorithms with one phase and without tagging is more preferable due to the fact that no additional transferring data through the network are needed. Approaches that sensitive to the data skew may be improved by optimizations with range partitioner. In case of data low selectivity semi-join algorithms can improve performance and reduce the possibility of memory overflow.

	Pre-processing	The number of phases	Tags	Sensitive to data skew	Need distr. cache	Memory overflow	Join algorithm
GRSJ	-	2	To value	yes	-	Number tuples for the same key is large	Nested loop
ORSJ	-	2	To key and value	yes	-	Number tuples for the same key is big	Nested loop
HYB	1 data	2	-	yes	-	Part size is large	Hash
MSPJ	2 data	1	-	yes	-	Part size is large	Hash
MSPMJ	2 data + sort	1	-	yes	-	-	Sort-merge
IMMJ	-	1	-	-	yes	Size of smaller dataset is large	Hash
MUL	1 data	1*part	-	-	yes	-	IMMJ

JDBM	-	1	-	-	yes	-	JDBM hash table
REV	-	1	-	-	yes	Part size is big and number of tuples for the same key is big	Hash

Table 1: Comparative analysis of algorithms.

The multiphase and JDBM map join algorithms is excluded from our experiments because of their poor performance.

4. Experiments

4.1. Dataset

Data are the set of tuples, which attributes are separated by a comma. Tuple is split into a pair of a key and a value, where value is the remaining attributes. Generation of synthetic data was done as in [4]. Join keys are distributed randomly except experiment with the data skew.

4.2. Cluster configuration

Cluster consists of three virtual machines, where one of them is master and slave at the same time, the remaining two are the slaves. Host configuration consists of 1 processor, 512 mb of memory for nodes, 5 gb is the disk size. Hadoop 20.203.0 runs on Ubuntu 10.10.

4.3. The General Case

The base idea of this experiment is to compare executions time of different phases of various algorithms. Some parameters are fixed: the number of Map and Reduce tasks is 3, the input size is $10^4 \times 10^5$ and $10^6 \times 10^6$ tuples.




Figure 1: Executions time of different phases of various algorithms. Size $10^4 * 10^5$.




Figure 2: Executions time of different phases of various algorithms. Size $10^6 * 10^6$.

For a small amount of data, Map phase, in which all tuples are tagged, and Shuffle phase, in which data are transferred from one phase to another, are more costly in Reduce-Side joins. It should be noted that GRSJ is better than ORSJ on small data, but it is the same on big data. It is because in first case time does not spend on combining tuples. Possible, on the larger data ORSJ outperform GRSJ when the usefulness of grouping by key will be more significant. Also for algorithms with pre-processing more time are spent on partitioning data. The algorithms in memory (IMMJ and REV) are similar in small data. Two algorithms are not shown in the

graph because of their bad times: JDBM-based map join and Multi-phase map join. In large data IMMJ algorithm could not be executed because of memory overflow.

4.4. Semi-Join

The main idea of this experiment is to compare different semi-join algorithms. These parameters are fixed: the number of Map and Reduce tasks is 3, the bitmap size of Bloom-filter is $25 * 10^5$, the number of hash-functions in Bloom-filter is 173, built-in Jenkins hash algorithm is used in Bloom-filter. Adaptive semi-join (ASGRSJ) does not finish because of memory overflow. The abbreviation of Bloom-filter semi-join for GRSJ is BGRSJ. The abbreviation of semi-join with selection for GRSJ is SGRSJ respectively.




Figure 3: Comparison of different semi-join implementations.

4.5. Speculative execution

Speculative execution reduces negative effects of non-uniform performance of physical nodes. In this experiment two join algorithms GRSJ and IMMJ is chosen because of different numbers of phases and one of them sensitive to the data skew. Two dataset are considered: normal data that consists of $10^5 * 10^5$ tuples and skew data that contain for one data $5 * 10^4$ same key in tuples and for second data 10 same keys in tuples. In case of IMMJ, which is not sensitive to the data skew, the performance with speculative execution is the similar approach without it. In case of GRSJ algorithm with uniform data approach without speculative execution is better than with it. But GRSJ algorithm with skew data and speculative execution outperforms four times approach without it.

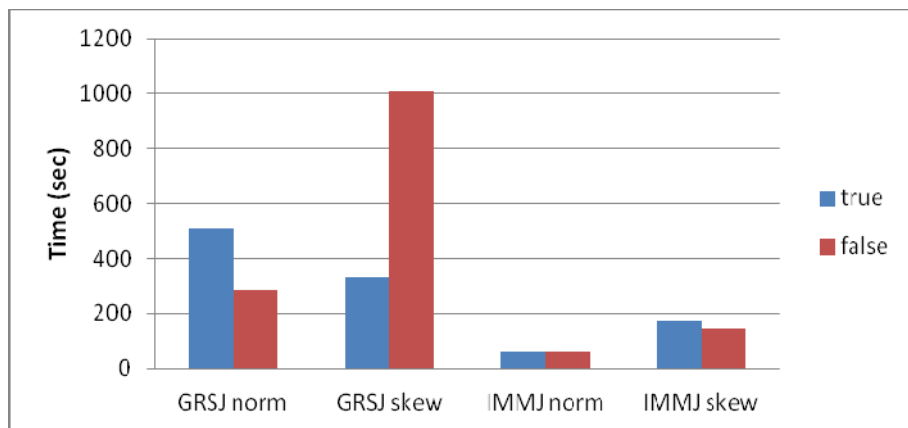


Figure 4: The effect of speculative execution.

4.6. Distributed cache

In [21] was showed that using of distributed cache is not always good strategy. They suggested that the problem can be a high speed network. This experiment was carried out for Reversed Map-Side join, because for which a distributed cache can be important. Replication was varied as 1, 2, 3 and size of data is fixed – $10^6 \times 10^6$ tuples. When data is small, the difference is not always visible. In large data algorithms with distributed cache outperform approach of reading from a globally distributed system.




Figure 5: Performance of Reversed Map-Side join with and without using distributed cache.

4.7. Skew data

It is known that many of the presented algorithms are sensitive to the data skew. In this experiment take part such algorithms as Reduce-side join with Simple Range-based Partitioner for GRSJ (GRSJRange) and Virtual Processor Partitioner for GRSJ (GRSJVirtual), and also for comparing in memory join: IMMJ, REV because of resistant to the skew. Fixed parameters are used: size of two dataset is 2×10^6 , one of the data set has the same key in 5×10^5 tuples, and another has the same keys in 10 or 1 tuples. In case with IMMJ was memory overflow.




Figure 6: Processing the data skew.

Although these experiments do not completely cover the tunable set of Hadoop parameters, they are shown the advantages and disadvantages of the proposed algorithms. The main problems of these algorithms are time spent on pre-processing, transferring data, the data skew, and memory overflow.

Each of the optimization techniques introduces additional cost to the implementation of the join, so the algorithm based on the tunable settings and specific data should be carefully chosen. Also important are the parameters of the network bandwidth when distributed cache are used or not used and a hardware specification of nodes because of it is importance when speculative executions are on. Speculative execution reduces negative effects of non-uniform performance of physical nodes.

Based on the collected statistics such as data size, how many keys will be taking part in the join, these statistics may be collected as well as the construction of a range partitioner, the query planner can choose an efficient variant of the join. For example, in [5] was proposed what-if analyses and cost-based optimization.

5. Future work

The algorithms discussed in this paper, only two sets are joined. It is interesting to extend from binary operation to multi argument joins. Among the proposed algorithms, there is no effective universal solution. Therefore, it is necessary to evaluate the proposed cost models for join algorithms. And for this problem it is need to use real cluster with more than three nodes in it and more powerful to process bigger data, due to the fact that the execution time on the virtual machine may be different from the real cluster in reading/writing, transferring data over the network and so on.

Also the idea of processing the data skew in MapReduce applications from [19] can be applied to the join algorithms. Another direction to future work is to extend algorithm to support a theta-join and outer join.

An interesting area for future work is to develop, implement and evaluate algorithms or extended algebraic operations suitable for complex similarity queries in an open distributed heterogeneous environment. The reasons to evaluate complex structured queries are: a need to combine search criteria for different types of information; a query refinement e.g. based on user profile or feedback; advanced users may need query structuring. The execution model and algebraic operation to be implemented are outlined in [31]. The main goal is to solve the problems presented in [8] as a problem.

In addition, one of the issues is efficient physical representation of data. Binary formats are known to outperform the text both in speed reading and partitioning key / value pairs, and the transmission of compressed data over the network. Along with the binary data format, column storage has already been proposed for paradigm MapReduce. It is interesting to find the best representation for specific data.

6. Conclusion

In this work we describe the state of the art in the area of massive parallel processing, presented our comparative study of these algorithms with optimizations such as semi-join and range partitioner. Also our directions of future work is discussed.

References

- [1] Daniel J. Abadi, Samuel R. Madden, and Nabil Hachem. Column-stores vs. row-stores: how different are they really? In Proceedings of the 2008 ACM SIGMOD international conference on Management of data, SIGMOD '08, pages 967–980, New York, NY, USA, 2008. ACM.
- [2] Azza Abouzeid, Kamil Bajda-Pawlikowski, Daniel Abadi, Avi Silberschatz, and Alexander Rasin. Hadoopdb: an architectural hybrid of mapreduce and dbms technologies for analytical workloads. Proc. VLDB Endow., 2:922–933, August 2009.
- [3] Foto N. Afrati and Jeffrey D. Ullman. Optimizing joins in a map-reduce environment. In Proceedings of the 13th International Conference on Extending Database Technology, EDBT '10, pages 99–110, New York, NY, USA, 2010. ACM.
- [4] Fariha Atta. Implementation and analysis of join algorithms to handle skew for the hadoop mapreduce framework. Master's thesis, MSc Informatics, School of Informatics, University of Edinburgh, 2010.
- [5] Shivnath Babu. Towards automatic optimization of mapreduce programs. In Proceedings of the 1st ACM symposium on Cloud computing, SoCC '10, pages 137–142, New York, NY, USA, 2010. ACM.
- [6] Spyros Blanas, Jignesh M. Patel, Vuk Ercegovic, Jun Rao, Eugene J. Shekita, and Yuanyuan Tian. A comparison of join algorithms for log processing in mapreduce. In Proceedings of the 2010 international conference on Management of data, SIGMOD '10, pages 975–986, New York, NY, USA, 2010. ACM.
- [7] A Chatzistergiou. Designing a parallel query engine over map/reduce. Master's thesis, MSc Informatics, School of Informatics, University of Edinburgh, 2010.
- [8] Surajit Chaudhuri, Raghu Ramakrishnan, and Gerhard Weikum. Integrating db and ir technologies: What is the sound of one hand clapping? In CIDR, pages 1–12, 2005.
- [9] Jeffrey Cohen, Brian Dolan, Mark Dunlap, Joseph M. Hellerstein, and Caleb Welton. Mad skills: new analysis practices for big data. Proc. VLDB Endow., 2:1481–1492, August 2009.
- [10] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: a flexible data processing tool. Commun. ACM, 53:72–77, January 2010.
- [11] Jeffrey Dean, Sanjay Ghemawat, and Google Inc. Mapreduce: simplified data processing on large clusters. In In OSDI04: Proceedings of the 6th conference on Symposium on Operating Systems Design & Implementation. USENIX Association, 2004.
- [12] Leonidas Fegaras, Chengkai Li, and Upa Gupta. An optimization framework for map-reduce queries. In EDBT 2012, march 2012.
- [13] Avriilia Floratou, Jignesh M. Patel, Eugene J. Shekita, and Sandeep Tata. Column-oriented storage techniques for mapreduce. Proc. VLDB Endow., 4:419–429, April 2011.
- [14] Alan F Gates. Programming Pig. O'Reilly Media, 2011.
- [15] Herodotos Herodotou. Hadoop performance models. CoRR, abs/1106.0940, 2011.
- [16] Herodotos Herodotou and Shivnath Babu. Profiling, what-if analysis, and cost-based optimization of mapreduce programs. PVLDB, 4(11):1111– 1122, 2011.
- [17] Eaman Jahani, Michael J. Cafarella, and Christopher R'e. Automatic optimization for mapreduce programs. Proc. VLDB Endow., 4:385–396, mar 2011.
- [18] Dawei Jiang, Anthony K. H. Tung, and Gang Chen. Map-join-reduce: Toward scalable and efficient data analysis on large clusters. IEEE Transactions on Knowledge and Data Engineering, 23:1299– 1311, 2011.
- [19] YongChul Kwon, Magdalena Balazinska, Bill Howe, and Jerome Rolia. A study of skew in mapreduce applications. Moscow, Russia, june 2011. In the 5th Open Cirrus Summit.
- [20] Yuting Lin, Divyakant Agrawal, Chun Chen, Beng Chin Ooi, and Sai Wu. Llama: leveraging columnar storage for scalable join processing in the mapreduce framework. In Proceedings of the 2011 international conference on Management of data, SIGMOD '11, pages 961–972, New York, NY, USA, 2011. ACM.
- [21] Gang Luo and Liang Dong. Adaptive join plan generation in hadoop. Technical report, Duke University, 2010.
- [22] Christine Morin and Gilles Muller, editors. European Conference on Computer Systems, Proceedings of the 5th European conference on Computer systems, EuroSys 2010, Paris, France, April 13-16, 2010. ACM, 2010.

- [23] Alper Okcan and Mirek Riedewald. Processing theta-joins using mapreduce. In Proceedings of the 2011 international conference on Management of data, SIGMOD '11, pages 949–960, New York, NY, USA, 2011. ACM.
- [24] Konstantina Palla. A comparative analysis of join algorithms using the hadoop map/reduce framework. Master's thesis, MSc Informatics, School of Informatics, University of Edinburgh, 2009.
- [25] Andrew Pavlo, Erik Paulson, Alexander Rasin, Daniel J. Abadi, David J. DeWitt, Samuel Madden, and Michael Stonebraker. A comparison of approaches to large-scale data analysis. In Proceedings of the 35th SIGMOD international conference on Management of data, SIGMOD '09, pages 165–178, New York, NY, USA, 2009. ACM.
- [26] Donovan A. Schneider and David J. DeWitt. A performance evaluation of four parallel join algorithms in a shared-nothing multiprocessor environment. SIGMOD Rec., 18:110–121, June 1989.
- [27] Rares Vernica, Michael J. Carey, and Chen Li. Efficient parallel set-similarity joins using mapreduce. In Proceedings of the 2010 international conference on Management of data, SIGMOD '10, pages 495–506, New York, NY, USA, 2010. ACM.
- [28] Vertica Systems, Inc. Managing Big Data with Hadoop & Vertica, 2009.
- [29] Guanying Wang, Ali Raza Butt, Prashant Pandey, and Karan Gupta. A simulation approach to evaluating design decisions in mapreduce setups. In MASCOTS, pages 1–11. IEEE, 2009.
- [30] Hung-chih Yang, Ali Dasdan, Ruey-Lung Hsiao, and D. Stott Parker. Map-reduce-merge: simplified relational data processing on large clusters. In Proceedings of the 2007 ACM SIGMOD international conference on Management of data, SIGMOD '07, pages 1029–1040, New York, NY, USA, 2007. ACM.
- [31] Anna Yarygina, Boris Novikov, and Natalia Vassilieva. Processing complex similarity queries: A systematic approach. In Maria Bielikova, Johann Eder, and A Min Tjoa, editors, ABDIS 2011 Research Communications: Proceedings II of the 5th East-European Conference on Advances in Databases and Information Systems 20 – 23 September 2011, Vienna, pages 212–221. Austrian Computer Society, September 2011.
- [32] Minqi Zhou, Rong Zhang, Dadan Zeng, Weining Qian, and Aoying Zhou. Join optimization in the mapreduce environment for column-wise data store. In Proceedings of the 2010 Sixth International Conference on Semantics, Knowledge and Grids, SKG '10, pages 97–104, Washington, DC.