

# Association Algorithm for Two Dynamically Enlarging Tables Implemented in Apache Spark

Alexander Agarkov  
JSC NICEVT  
Moscow, Russia  
a.agarkov@nicevt.ru

Alexander Semenov  
JSC NICEVT  
Moscow, Russia  
alxdr.semenov@gmail.com

## Abstract

In the paper we consider association problem with constraints for two dynamically enlarging tables. We consider an ordered set of rule groups which determine associations between entries from the first table and the second table. Each entry is associated with other entries from both tables directly or indirectly through the other associations. In the problem it is needed to list the associated entries for each entry. Tables are dynamically enlarging, the goal is to improve potential performance of the association process by using of the previously built associations. We consider a base full association algorithm and propose a partial association algorithm that improves the efficiency of the base algorithm, implement and evaluate both algorithms in Apache Spark for a particular case on 12 cluster nodes.

**Keywords:** association problem, dynamically enlarging tables, bipartite dynamic graph

## 1 Introduction

In the recent years data intensive applications have become widespread and appeared in many science and engineering areas (biology, bioinformatics, medicine, cosmology, finance, social network analysis, cryptology etc.). They are characterized by a large amount of data, irregular workloads, unbalanced computations and low sustained performance of computing systems. Development of new algorithmic approaches and programming technologies are urgently needed to boost

*Copyright © by the paper's authors. Copying permitted for private and academic purposes.*

In: V. Voevodin, A. Simonov (eds.): Proceedings of the GraphHPC-2017 Conference, Moscow State University, Russia, 02-03-2017, published at <http://ceur-ws.org>.

efficiency of HPC systems for similar applications, thus enabling advancing of HPC and Big Data convergence [1].

In the paper we consider an association problem for two dynamically enlarging tables. We have two large tables and an ordered set of rule groups which determine associations between entries from the first table and the second table. When two table entries compose an association by a rule in the current rule group, then these entries must be excluded from association process for the following rule groups. Each entry is associated with other entries from both tables directly or indirectly through the other associations. It is needed to determine the association type and to list the associated entries for each entry. Tables are dynamically enlarging, the goal is to improve potential performance of the association process by using of the associations, built on the original tables.

Spark [2] is a framework which optimizes programming and execution models of MapReduce [3] by introducing resilient distributed dataset (RDD) abstraction. Users can choose between the cost of storing an RDD, the speed of accessing it, the probability of losing part of it, and the cost of recomputing it. Apache Spark [4] is a popular open-source implementation of Spark.

In the paper we describe a base full association approach to the problem, propose a partial association approach that improves the efficiency of the base approach, implement corresponding algorithms using Apache Spark and present evaluation results on a cluster.

## 2 Association Problem for Two Dynamically Enlarging Tables

We consider two large *tables* with  $M$  rows. The tables have identical structure, each table has  $N$  fields. The table unique key consists of all  $N$  fields of the table.

We consider an ordered set of rule groups, which determine associations between entries from the first table and entries from the second table:

- *Rule* is a set of fields, which are used to compare two table entries. It is required to build associations between the tables: to find matches between different table entries by the rule.
- *Group* is a set of rules; rules of a group are applied to the table entries independently of each other. When two table entries compose the association by a rule in the current rule group, then these entries are marked by the current group number and must be excluded from association process for the following rule groups.

Each table entry can be associated with one or many entries of another table. Moreover, association is a transitive relation. Associations for each entry can be classified into one of four association types: one from the first table to one from the second table, one to many, many to one and many to many (1: 1, 1: M, M: 1, M: N). Therefore each entry is associated with other entries from the both tables directly or indirectly through the other associations.

The goal of the **table association problem** is to determine the association type and to list of the associated entries for each entry.

After we build associations between the tables,  $K$  new entries are added to each table. Added entries differ from original entries by a given subset of fields. Association process is needed to repeat to make the augmented tables associated too.

Full association approach can generate associations between given tables by the mentioned set of rules from scratch. It is required to build associations between the augmented tables.

The goal of the **dynamically enlarging table association problem** is to improve potential performance of the full association approach on the augmented tables by using of the associations, built on the original tables.

For the sake of simplicity in the paper we consider a particular case of the problem.

## 2.1 Data Structure and Association Rules

In our work each table entry has 5 fields, where the first and the second fields are integer identifiers, three other fields are data fields. The unique key for every entry is all of the five fields. Each entry has a unique synthetic identifier.

In the work the considered ordered set of rule groups consists of 5 groups and 15 rules, see Table 1. Symbol "+" denotes equality requirement of the corresponding fields of the two table entries. Symbol "-" denotes that

fields of the two table entries are not matched. For example, a key that determines an association between two table entries is specified for each rule and consists of the fields, that are marked with the "+" symbol.

group	rule	ID 1	ID 2	#1	#2	#3
1	1	+	+	+	+	+
2	2	+	+	+	+	-
2	3	+	+	+	-	+
2	4	+	+	-	+	+
3	5	+	+	-	+	-
3	6	+	+	+	-	-
3	7	+	+	-	-	+
4	8	+	-	+	+	+
4	9	+	-	+	+	-
4	10	+	-	+	-	+
4	11	+	-	-	+	+
5	12	-	+	+	+	+
5	13	-	+	+	+	-
5	14	-	+	+	-	+
5	15	-	+	-	+	+

Table 1: The description of the considered rules groups.

The full association approach is matching each entry from the first table with each entry of the second table by each rule from the current rule group. If matching is successful then we create and store an association between the entries; the association is marked by the current group number. Entries that do not have any associations are marked by the group number six.

## 3 Algorithms

### 3.1 Full Association Algorithm

The full association algorithm consists of two stages: associations matching and transitive closure. The first stage actually implements the full association approach. All possible pairs are found by the first group of rules, then every entry that is included in the pairs is excluded from the tables. This procedure is repeated for each group of rules. The result of the stage is a set of associations (future graph edges) between entries (future graph vertices).

At the second stage the transitive closure (TC) algorithm is executed for each selected group. At first, we construct a bipartite graph. Vertices in the left vertex set are unique identifiers of the entries from the first table, in the right vertex set there are unique identifiers of the entries from the second table. There exists an edge between two vertices of the different graph parts if the association between corresponding entries has been found during the first stage of the algorithm.

Transitive closure [5]  $TC$  is performed by the following formula:

$$TC = \cup_{i=1,2..R_i}, \text{ where } R_{i+1} = R_i \text{ join } E, \\ R_1 = E, E - \text{set of graph edges.} \quad (1)$$

The transitive closure is built by repetitive merging result of a join operation between previous resulting set of edges and original set of graph edges until the result is not changed, i.e. fixed point is reached. Thus, for each vertex in  $TC$  there exist vertex pairs that connect current vertex with other vertices in the connected component.

Potentially it is possible to use better algorithm for implementing transitive closure, e.g. [6].

Finally, the association type of each vertex is defined (1:1, 1:M, M:1, M:N).

### 3.2 Partial Association Algorithm

There are original (old) tables, the associations that have been built for old tables, and there are new tables that are smaller than original. The added entries differ from original entries by the #3 field.

When new entries are added to the original tables, one can apply full association algorithm to the augmented tables from scratch. We propose partial association algorithm to improve performance by using of associations that are built for the original tables.

The partial association algorithm is executed also in two stages: association matching and transitive closure.

It is important that added (new) entries differ from the original entries by the #3 field. The main idea of first stage is matching only new entries of the tables for each rule with matching requirement by the #3 field. In the case there will be no associations between new and old entries. For each rule without matching requirement by the last field new and old entries must be matched, see Figure 1.




Figure 1: Partial association algorithm. Process a. with b. without #3 field matching requirement.

The association matching stage slightly differs from the same stage in the full association algorithm. Each entry included in new associations must be excluded

from old associations. As seen in Figure 2, if a new entry is associated with an old entry, and the association group number  $new\_gn$  is smaller than the group number  $old\_gn$  of the association between the old entry with another entry, then these old associations must be removed; if  $new\_gn$  is equal to  $old\_gn$  then the new entry should be added to the component.

The transitive closure stage is executed only for new associations. The resulting graph of transitive closure is combined with the old graph with invalid associations excluded during the matching stage.

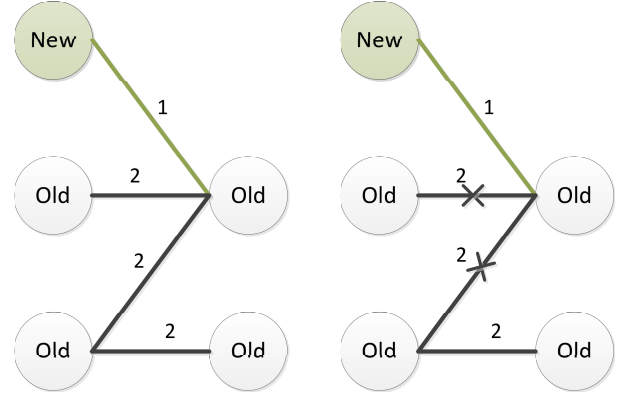


Figure 2: Association of new entry with old entry and removing of old associations.

## 4 Implementation Details

Apache Spark [4] is a popular open-source implementation of Spark. It provides programmers with an application programming interface centered on a data structure called the resilient distributed dataset (RDD), a read-only multiset of data items distributed over a cluster of machines, that is maintained in a fault-tolerant way [2]. It was developed in response to limitations in the MapReduce cluster computing paradigm [3], which forces a particular linear dataflow structure on distributed programs: MapReduce programs read input data from disk, map a function across the data, reduce the results of the map, and store reduction results on disk. Spark's RDDs function as a working set for distributed programs that offers a (deliberately) restricted form of distributed shared memory [7]. The latest Spark program interface DataFrame [8] seems to be more efficient than RDD, but in the current work we use RDD, and we suppose to use DataFrame in the next research works.

We use Java 8 and Apache Spark 1.6.1 for implementation of the full and partial association algorithms. We use RDD of `Tuple5<Long, Long, Long, Long, Long>` type for table structure representation, the sequence of types in `Tuple5` corresponds to the table fields ID1, ID2, #1, #2, #3. We attach unique

rule	% in tables
1	72
2	2
3	6
4	12
5	0.1
6	0.5
7	0.4
8	0.5
9	0.01
10	0.04
11	0.35
12	0.83
13	0.02
14	0.12
15	0.26
no associations	4.86

Table 2: Distribution of the second table modifications in the rules implemented in the synthetic table generator.

identifier (**Long**) to the **Tuple5** of each entry.

After the association stage we have RDD of **Tuple2<Long, Long>** that represents association between the unique identifiers of two table entries.

#### 4.1 Synthetic Table Generator

Synthetic table generator creates distributed random tables and works as follows. Firstly, two identical tables of the required size are generated. Each field value is a uniformly distributed random integer number in the following intervals:

- ID1, ID2 – [0; 10000),
- #1 – [0; 1000000),
- #3 – [0; 1000).

Value of the #2 field is a position number of the entry.

We randomly modify second table entries in order to create possibility of association between entries from the first and the second tables for each rule. We modify entry fields that are marked with the “\_” symbol in Table 1. Distribution of modifications in the rules is shown in Table 2. 72% of the second table entries remain unchanged. In 2% of the table entries there are random modifications in the #3 field values. In 6% of the table entries there are random modifications in the #2 field values and so on. As a result 72% of the table entries correspond to the first rule, 2% – to the second rule, 6% – to the third rule and so on.

We generate the augmented tables as follows. Firstly, we add new entries to the first table, field values of each entry is a uniformly distributed random integer number in the following intervals:

- ID1, ID2 – [0; 10000),
- #1 – [0; 1000000),
- #3 – [1000; 2000).

Value of the #2 field is a position number of the new entry in the whole table. As can be seen the old and the new tables have different values in the #3 field.

Secondly, we copy augmented part of the first table to the second table and randomly modify it as well as is described in Table 2.

## 5 Performance Evaluation

All presented results are obtained on the Angara-K1 cluster. We use 12 out of 36 nodes, each of the 12 nodes is equipped with a 8-core Intel Sandy-Bridge Xeon E5-2660 processor (2.2 GHz) and 64 GB DDR3 DRAM. All Angara-K1 nodes are linked by the Angara interconnect and 1 Gbit/s Ethernet. High-speed Angara interconnect is developed in NICEVT, performance evaluation of the Angara-K1 cluster with Angara interconnect on scientific workloads is presented in [9].

Figures 3, 4 and 5 show the comparison results of the full and partial association algorithms. The reported running time does not include reading data and writing the result to the filesystem.

The algorithm running times are shown in Figure 3, we use 8 cores per node and 8 nodes of the cluster, table size is varied. Old table size is 300 million entries, new table size is 75 million entries. The figure shows that performance difference between the algorithms grows with table size.

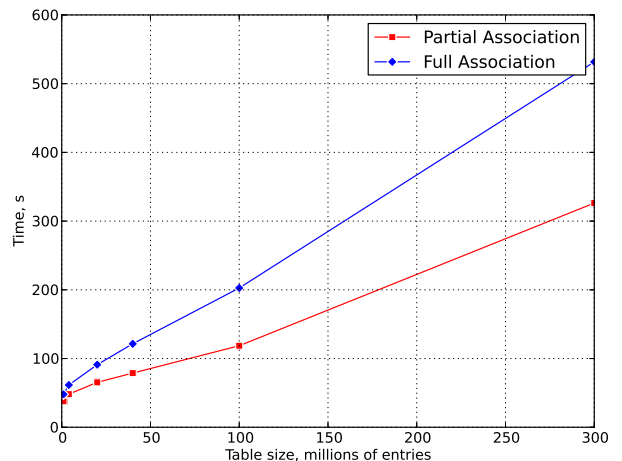


Figure 3: Running times of the full and partial association algorithms on the different table size. 8 cores per each of the 8 Angara-K1 cluster nodes.

Strong scaling is shown in Figure 4. Old table size is 100 million entries, new table size is 25 million entries. The speedup of the full and partial association

algorithms is approximately 3 on 8 nodes. Among the possible reasons of moderate performance there is a single one: Spark configuration is not optimal. Further tuning can address the problem. Horizontal line from 8 to 12 nodes indicates that table size is small for further performance increasing.

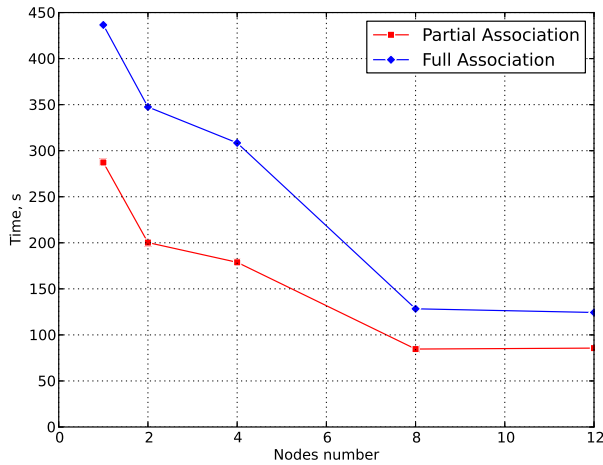


Figure 4: Strong scaling results for the old tables with 100 million entries and new tables with 25 million entries.

The algorithm running times on the amount of new data are shown in Figure 5. We use 6 nodes, 300 million entries in the each table, fraction of the new table entries varies from 12.5 to 100 percents of the total table size. The smaller the percentage of new data is, the faster the partial association algorithm is executed. The running time of the full association algorithm does not change, because the total amount of data does not change.




Figure 5: Algorithm performance comparison for a various ratio of the new and the total table sizes.

## 6 Conclusion

In the paper we propose the partial association algorithm for the table association problem of two dynamically enlarging tables with specific constraints. For the sake of simplicity we consider a particular case of the problem. We implement the base full association algorithm and the proposed algorithm using Apache Spark and present performance evaluation of the algorithms on the cluster Angara-K1. Performance of the proposed algorithm exceeds performance of the full association algorithm for a variety of data sets.

The work was supported by the grant No. 17-07-01592A of the Russian Foundation for Basic Research (RFBR).

## References

- [1] D. Reed and J. Dongarra, “Exascale computing and Big Data: The next frontier,” *Communications of the ACM*, vol. 57, no. 7, pp. 56–68, 2014. <http://www.netlib.org/utk/people/JackDongarra/PAPERS/Exascale-Reed-Dongarra.pdf> (accessed: 11.10.2017).
- [2] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica, “Spark: Cluster computing with working sets,” *HotCloud*, vol. 10, p. 7, 2010. [https://www.usenix.org/legacy/event/hotcloud10/tech/full\\_papers/Zaharia.pdf](https://www.usenix.org/legacy/event/hotcloud10/tech/full_papers/Zaharia.pdf) (accessed: 11.10.2017).
- [3] J. Dean and S. Ghemawat, “MapReduce: Simplified data processing on large clusters,” in *Proceedings of the 6th Conference on Symposium on Operating Systems Design and Implementation - Volume 6*, OSDI’04, (Berkeley, CA, USA), USENIX Association, 2004. <https://static.googleusercontent.com/media/research.google.com/en//archive/mapreduce-osdi04.pdf> (accessed: 11.10.2017).
- [4] *Apach Spark Homepage*. <http://spark.apache.org/> (accessed: 11.10.2017).
- [5] S. Abiteboul, R. Hull, and V. Vianu, eds., *Foundations of Databases: The Logical Level*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 1st ed., 1995. <http://webdam.inria.fr/Alice/pdfs/all.pdf> (accessed: 13.10.2017).
- [6] E. Nuutila, “Efficient transitive closure computation in large digraphs, mathematics and computing in engineering series no. 74 phd thesis helsinki university of technology,” 1995. <http://www.cs.hut.fi/~enu/thesis.pdf> (accessed: 13.10.2017).

- [7] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. J. Franklin, S. Shenker, and I. Stoica, “Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing,” in *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation*, NSDI’12, (Berkeley, CA, USA), USENIX Association, 2012. [http://www-bcf.usc.edu/~minlanyu/teach/csci599-fall112/papers/nsdi\\_spark.pdf](http://www-bcf.usc.edu/~minlanyu/teach/csci599-fall112/papers/nsdi_spark.pdf) (accessed: 11.10.2017).
- [8] M. Armbrust, R. S. Xin, C. Lian, Y. Huai, D. Liu, J. K. Bradley, X. Meng, T. Kaftan, M. J. Franklin, A. Ghodsi, *et al.*, “Spark SQL: Relational data processing in Spark,” in *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, pp. 1383–1394, ACM, 2015. <https://amplab.cs.berkeley.edu/wp-content/uploads/2015/03/SparkSQLSigmod2015.pdf> (accessed: 11.10.2017).
- [9] A. Agarkov, T. Ismagilov, D. Makagon, A. Semenov, and A. Simonov, “Performance evaluation of the Angara interconnect,” in *Proceedings of the International Conference Russian Supercomputing Days*, pp. 626–639, 2016. <http://www.dislab.org/docs/rsd2016-angara-bench.pdf> (accessed: 11.10.2017).