

Efficient FP Growth using Hadoop - (Improved Parallel FP-Growth)

Sankalp Mitra¹, Suchit Bande², Shreyas Kudale³, Advait Kulkarni⁴, Asst. Prof. Leena A. Deshpande⁵

^{1,2,3,4} Department of Computer Engineering, VIIT, Pune

⁵ Assistant Professor, Department of Computer Engineering, VIIT, Pune

Abstract- As an important part of discovering association rules, frequent itemsets mining plays a key role in mining associations, correlations, causality and other important data mining tasks. Since some traditional frequent itemsets mining algorithms are unable to handle massive small files datasets effectively, such as high memory cost, high I/O overhead, and low computing performance, an improved Parallel FP-Growth (IPFP) algorithm and discuss its applications in this paper. In particular, a small files processing strategy for massive small files datasets to compensate defects of low read/write speed and low processing efficiency in Hadoop. Moreover, use of MapReduce to implement the parallelization of FP-Growth algorithm, thereby improving the overall performance of frequent itemsets mining. The experimental results show that the IPFP algorithm is feasible and valid with a good speedup and a higher mining efficiency, and can meet the rapidly growing needs of frequent itemsets mining for massive small files datasets.

Index Terms- Frequent itemsets mining, Hadoop MapReduce, Parallel FP-Growth, Small files problem

I. INTRODUCTION

As data management and analysis are facing new challenges in the age of big data, the rationality and timeliness of the data processing methods are becoming the research hotspot of big data statistical analysis, and big data association analysis greatly increases the profits of enterprises.

As one of the important research directions of data mining, frequent itemsets mining plays an essential role in mining associations, correlations, causality and other important data mining tasks [1] which is a strong impetus to the applications of association rules in markets selection, decision analysis and business management [5]. The existing classical frequent itemsets mining algorithms are breadth-first algorithm Apriori proposed by Agrawal et al. in 1994 and depth-priority algorithm FP-Growth presented by Han et al. in 2000. In order to improve the efficiency of frequent itemsets mining, many researchers have proposed several methods to optimize classical algorithms [11]; [10]. Meanwhile, to address the bottleneck of mining performance and reduce memory consumption and computation cost of the machine in the single machine environment, parallel and distributed algorithms are proposed methods [2], the applications of corresponding algorithms for frequent itemsets mining in the large-scale datasets will easily cause high CPU consumption, high memory cost, high I/O overhead, low computing performance and other issues.

As typical methods of Hadoop Distributed File System (HDFS) and MapReduce parallel programming model provide a new idea for handling big data. In the frequent itemsets mining for large-scale data, a MapReduce approach of parallel FP-Growth (PFP) algorithm is proposed in [8], and the performance of PFP algorithm is enhanced by adding load balancing features in [20], but these methods ignore frequent

Itemsets mining for massive small files datasets in Hadoop.

With the arrival of big data era, massive data are growing rapidly. However, in reality, most of the large-scale data are composed of massive small files. Small files usually refer to those file sizes, which are less than 64 MB. According to a study in 2007 at the National Energy Research Scientific Computing Center, 43% of the over 13 million files on a shared parallel file system are under 64 KB and 99% are under 64 MB (Petascale Data Storage

Institute (2007)), and more scientific applications consist of a large number of small files are depicted in [4]. Nevertheless, in the face of massive small files datasets, the constructed FP-tree in Parallel FP-Growth (PFP) algorithm cannot fit into the memory, which often causes problems such as memory overflow and huge communication overhead. Meanwhile, the computing efficiency of the Hadoop platform largely depends on the performance of HDFS and MapReduce [17], and Hadoop was, at first, designed specifically to handle streaming large files, so when dealing with massive small files, there are significant limitations. Massive small files will reduce the performance of Hadoop, which is mainly shown in the following two aspects:

- (1) The access efficiency of HDFS is decreased.
- (2) The additional overhead of MapReduce is increased. [21]

II. IPFP ALGORITHM DESCRIPTION

IPFP algorithm for mining frequent itemsets in massive small files datasets in detail.

(1) Write a small files processing program—*Sequence File*. The *Sequence File* is used to merge all massive small files, which are composed of a large amount of transaction datasets stored in HDFS, into a large transaction data file (transaction database).[21]

(2) Equally divide the transaction database into several sub-transaction databases and then assign them to different nodes in Hadoop cluster. This step is automatically operated by HDFS, when necessary we can use the balance command enabling its file system to achieve load balancing. [21]

(3) Compute *support* count of each item in the transaction database by MapReduce, and then obtain the set of *I_list* from *support* count in descending order.[17]

(4) Divide *I_list* into *M* groups, denoted as *Group_list* (abbreviated as *G_list*), and assign *group_id* for each group sequentially and each *G_list* contains a set of items.[3]

(5) Complete the parallel computing of FP-Growth algorithm by MapReduce. The *Map* function compares the item of each transaction in the sub-transaction database with the item in *G_list*. If they are same, then distribute the corresponding transaction to the machine associated with *G_list*. Otherwise, compares with the next item in *G_list*. Eventually, the independent sub-transaction databases corresponded to *G_list* will be produced. The *Reduce* function recursively computes the independent sub-transaction databases generated in step , and then constructs the FP-tree. This step is similar to the process of traditional FP-tree generation, but the difference is a size *K* max-heap *HP* which stores frequent pattern of each item.[21]

(6) Aggregate the local frequent itemsets generated from each node in the cluster by MapReduce, and finally get the global frequent itemsets.

III. IPFP ALGORITHM IMPLEMENTATION STEPS

Implement the IPFP algorithm, which is mainly composed of four steps as described in the following.[21]

Step 1: Merge massive small files. *Sequence File* is composed of a series of <key, value>, where the key is the name of small files and the value is the content of small files before merging. *Sequence File* exploits three classes—*WholeFileInputFormat*, *WholeFileRecordReader* and *SmallFilesToSequenceFileConverter*, to merge the massive small files into a large file.

(1) *WholeFileInputFormat* class:

1) The *isSplittable* () method reloads and returns the value of false, and the purpose is to maintain the input file not to be divided into slices.

2) The *getRecordReader* () method returns a customized *RecordReader*.

(2) *WholeFileRecordReader* class: the *FileSplit* is converted into a record, where the *key* of the record is the filename and the *value* is the content of this file.

(3) *SmallFilesToSequenceFileConverter* class: Massive small files are merged into a sequential file, and this class contains the *Map* () and the *Reduce* (). The input format of data is *WholeFileInputFormat*, while the output format is *SequenceFileOutputFormat*.

Step 2: Compute *I_list*. The complexity of time and space is $O(TDBsize/P)$, (*TDBsize*: the size of transaction database, *P*: the number of parallel MapReduce programs).

Step 3: Generate *G_list* from *I_list*, and complete parallel computing of FP-Growth algorithm. *Map* () judges which *G_list* the item of transactions in the machine belongs to, and then sends this transaction to the corresponding *G_list* machine. In order to avoid sending duplicate transaction, we delete the duplicate entries in the hash table. Each *Reduce* () handles the independent sub-transaction database associated with *G_list*, which creates heap *HP* with a size of *K* for each item in *G_list*.

Step 4: Aggregate local frequent itemsets generated in the previous step from each node, and then we get global frequent itemsets.

IV. EXPERIMENTS AND RESULTS

In this section, we show the study results from different sources that the feasibility, speedup, validity and efficiency of IPFP algorithm by two experiments. The experimental Hadoop cluster is composed of one Master machine and four Slave machines with Intel Pentium® Dual-Core E5700 3.00GHz CPU and 2.00GB RAM.

All the experiments are performed on Ubuntu 12.04 OS with Hadoop 0.20.2, Jdk 1.6.0 and Eclipse 3.7.1. The real data from the Frequent Itemset Mining Dataset Repository are used as the experimental data, which are processed into three groups of different sizes datasets. The feasibility, validity, speedup and efficiency are used to evaluate the overall performance of IPFP algorithm and compare it with the PFP algorithm in the same environment.[21]

V. APPLICATIONS

In this section, we apply the IPFP algorithm to a car database which consist different attributes of car such as, car model number, mileage, etc. The IPFP algorithm is used to analyze the difference between the time taken to run a regular fp growth to IPFP algorithm.

The IPFP algorithm shows better processing performance and a higher mining efficiency than PFP algorithm.

VI. CONCLUSIONS

In this paper, it is described that the small files processing strategy, the IPFP algorithm can reduce memory cost greatly and improve the efficiency of data access, thus avoids memory overflow and reduces I/O overhead. Meanwhile, the IPFP algorithm is migrated to the MapReduce environment, which can complete frequent itemsets mining efficiently and thus enhance the overall performance of FP-Growth algorithm. The experimental results show that IPFP algorithm can make a breakthrough where PFP algorithm has its defects in handling massive small files datasets, and has a good speedup and a higher mining efficiency.

REFERENCES

- [1] Agrawal, R. and Shafer, J. C. (1996) "Parallel mining of association rules," IEEE Transactions on Knowledge and Data Engineering, 8(6), 962-969.
- [2] Agrawal, R. and Srikant, R. (1994) "Fast algorithms for mining association rules," In: Proceedings of the 1994 international conference on very large data bases (VLDB' 94), Santiago, Chile, 487-499.
- [3] Brin, S., Motwani, R. and Silverstein, C. (1997) "Beyond market basket: generalizing association rules to correlations," In: Proceeding of the 1997 ACM SIGMOD international conference on management of data (SIGMOD'97), Tucson, AZ, 265-276.
- [4] Carns, P., Lang, S., Ross, R., Vilayannur, M., Kunkel, J. and Ludwig, T. (2009) "Small File access in parallel file systems," In: Proceeding of the 2009 IEEE international symposium on parallel and distributed processing, Rome, Italy, 1-11.

- [5] Cil, I. (2012) "Consumption universes based supermarket layout through association rule mining and multidimensional scaling," *Expert Systems with Applications*, 39(10), 8611-8625.
- [6] Grahne, G., Lakshmanan, L. and Wang, X. (2000) "Efficient mining of constrained correlated sets," In: *Proceeding of the 2000 international conference on data engineering (ICDE'00)*, San Diego, CA, 512-521.
- [7] Han, J., Pei, J. and Yin, Y. (2000) "Mining frequent patterns without candidate generation," In: *Proceeding of the 2000 ACM SIGMOD international conference on management of data (SIGMOD'00)*, Dallas, TX, 1-12.
- [8] Li, H., Wang, Y., Zhang, D., Zhang, M. and Chang, E. Y. (2008) "PPF: Parallel FP-Growth for query recommendation," In: *Proceeding of the 2008 ACM conference on Recommender systems*, Lausanne, Switzerland, 107-114.
- [9] Liao, S., Chen, Y. J. and Deng, M. (2010) "Mining customer knowledge for tourism new product development and customer relationship management," *Expert Systems with Applications*, 37(6), 4212-4223.
- [10] Lin, K. C., Liao, I. E. and Chen, Z. S. (2011) "An improved frequent pattern growth method for mining association rules," *Expert Systems with Applications*, 38(5), 5154-5161.
- [11] Park, J. S., Chen, M. S. and Yu, P. S. (1995) "An effective hash-based algorithm for mining association rules," In: *Proceeding of the 1995 ACM-SIGMOD international conference on management of data (SIGMOD'95)*, San Jose, CA, 175-186.
- [12] Petascale Data Storage Institute. (2007) "NERSC file system statistics," World Wide Web electronic publication, Online, Available: <http://pdsi.nersc.gov/filesystem.htm>.
- [13] Pramudiono, I. and Kitsuregawa, M. (2003) "Parallel FP-Growth on PC Cluster," *Advances in Knowledge Discovery and Data Mining*, 2637, 467-473.
- [14] Silverstein, C., Brin, S., Motwani, R. and Ullman, J. (2000) "Scalable techniques for mining causal structures," *Data Mining and Knowledge Discovery*, 4(2-3), 163-192.
- [15] Tanbeer, S. K., Ahmed, C. F. and Jeong, B. S. (2009) "Parallel and distributed algorithms for frequent pattern mining in large databases," *IETE Technical Review*, 26(1), 55-65.
- [16] Tsay, Y. J., Hsu, T. J. and Yu, J. R. (2009) "HUT: A new method for mining frequent itemsets," *Information Sciences*, 179(11), 1724-1737.
- [17] White, T. (2012) *Hadoop: The Definitive Guide*, 3rd ed., O'Reilly Media Inc, Sebastopol, CA.
- [18] Zaïane, O. R., El-Hajj, M. and Lu, P. (2001) "Fast parallel association rules mining without candidacy generation," In: *Proceeding of the 2001 international conference on data mining (ICDM'01)*, San Jose, CA, 665-668.
- [19] Zaki, M. J., Parthasarathy, S., Ogihara, M. and Li, W. (1997) "Parallel algorithm for discovery of association rules," *Data Mining and Knowledge Discovery*, 1(4), 343-373.
- [20] Zhou, L., Zhong, Z., Chang, J., Li, J., Huang, J. Z. and Feng, S. (2010) "Balanced parallel FP-Growth with MapReduce," In: *Proceeding of the 2010 IEEE Youth conference on information computing and telecommunications (YC-ICT'10)*, Beijing, China, 243-246.
- [21] Dawen Xia, Yanhui Zhou, Zhuobo Rong and Zili Zhang "IPFP: An Improved Parallel FP-Growth Algorithm for Frequent Itemsets Mining", School of Computer and Information Science, Southwest University, Chongqing, China, Institute of Statistics, Southwest University, Chongqing, China, School of Information Technology, Deakin University, Victoria, Australia Guizhou Minzu University, Guiyang, China, Corresponding author: Zili Zhang,

AUTHORS

First Author – Sankalp Mitra, Department of Computer Engineering, VIIT, Pune
Second Author – Suchit Bande, Department of Computer Engineering, VIIT, Pune
Third Author – Shreyas Kudale, Department of Computer Engineering, VIIT, Pune
Fourth Author – Advait Kulkarni, Department of Computer Engineering, VIIT, Pune
Fifth Author – Asst. Prof. Leena A. Deshpande, Assistant Professor, Department of Computer Engineering, VIIT, Pune