



Big Data Analysis using HDFS, C-MEANS and MapReduce

¹Shalini Jain, ²Satendra Sonare, ³Ashok Verma

¹Research Scholar, ²Asst. Prof., ³Prof. & Head of Deptt.

^{1,2,3} Department of Computer Science & Engineering,
Gyan Ganga Institute of Technology & Sciences,
Jabalpur (M.P.), India

Abstract— In this electronic age, increasing number of organizations are facing the problem of explosion of data and the size of the databases used in today's enterprises has been growing at exponential rates. Data is generated through many sources like business processes, transactions, social networking sites, web servers, etc. and remains in structured as well as unstructured form [2]. The term "Big data" is used for large data sets whose size is beyond the ability of commonly used software tools to capture, manage, and process the data within a tolerable elapsed time. Big data sizes are a constantly moving target currently ranging from a few dozen terabytes to many peta bytes of data in a single data set. Difficulties include capture, storage, search, sharing, analytics and visualizing. Big data is available in structured, unstructured and semi-structured data format. Relational database has fails to store this multi-structured data. Apache Hadoop is efficient, robust, reliable and scalable framework to store, process, transforms and extract big data. Hadoop framework is open source and free software which is available at Apache Software Foundation. In this paper we will present Hadoop, HDFS, Map Reduce and c-means big data algorithm to minimize efforts of big data analysis using Map Reduce code. The objective of this paper is to summarize the state-of-the-art efforts in clinical big data analytics and highlight what might be needed to enhance the outcomes of clinical big data analytics tools and related fields.

Index Terms—Big data mining, heterogeneity, HDFS, Map Reduce, hadoop

I. INTRODUCTION

Big data is the term used to describe huge datasets having the "3 V" definition: volume, variety, velocity and value (e.g. medical images, electronic medical records (EMR), biometrics data, etc.). Such datasets present problems with storage, analysis, and visualization [1,2]. To deal with these challenges, new software programming frameworks to multithread computing tasks have been developed [2-4].

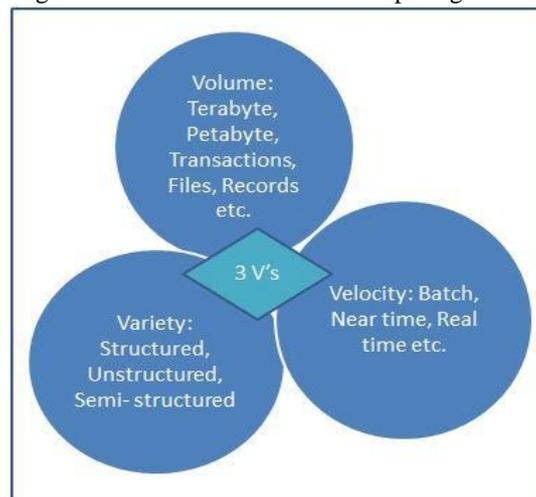


Fig. 1 Three V's of Big Data

These programming frameworks are designed to get their parallelism not from a supercomputer, but from computing clusters: large collections of commodity hardware, including conventional processors (computing nodes) connected by Ethernet cables or inexpensive switches. These software programming frameworks begin with a new form of file system, known as a distributed file system (DFS) [3,4], which features much larger units than the disk blocks in a conventional operating system. DFS also provides replication of data or redundancy to protect against the frequent media failures that occur when data is distributed over potentially thousands of low cost computing nodes [3]. The goal of this review is to summarize the potential and expanding usage of

MapReduce on top of the Hadoop platform in the processing of clinical big data. A secondary objective is to highlight the potential benefits of predictive and prescriptive clinical big data analytics. These types of analytics are needed for better usage and optimization of resources [5,6]. Types of analytics Analytics is a term used to describe various goals and techniques of processing a dataset.

There are three types of analytics:

- 1- Descriptive analytics: is a process to summarize the dataset under investigation. It may be used to generate standard reports that might be useful to address questions like "What happened? What is the problem? What actions are needed?"
- 2- Predictive analytics: descriptive analytics, unfortunately do not tell anything about the future, that is the reason predictive analytics is needed. Predictive analytics utilize statistical models of the historical datasets to predict the future. The predictive ability is dependent on the goodness of fit of the statistical model [6].
- 3- Prescriptive analytics: are the type of analytics that help in utilizing different scenarios of the data model (i.e. multi-variables simulation, detecting hidden relationships between different variables. Prescriptive analytics are generally used in optimization problems and require sophisticated algorithms to find the optimum solution and therefore are less widely used in some fields (i.e. clinical big data analytics).

II. HIGH PERFORMANCE COMPUTING (HPC) SYSTEMS

Distributed system

A distributed system [3] is a setup in which several independent computers (computing nodes) participate in solving the problem of processing a large volume of and variety of structured/semi-structured/unstructured data.

Grid computing system

The grid computing system [7] is a way to utilize resources (e.g. CPUs, storage of computer systems across a worldwide network, etc.) to function as a flexible, pervasive, and inexpensive accessible pool of computing resources that can be used on demand by any task. Furthermore the system architecture of GPUs may not be suitable for the MapReduce architecture and may require a great deal of modification [9]. The basic differences between grid computing and distributed computing systems are:

1. A distributed computing system manages hundreds or thousands of computer systems, which are limited in processing resources (e.g. memory, CPU, storage, etc.). However the grid computing system is concerned about efficient usage of heterogeneous systems with optimal workload management servers, networks, storage, etc.
2. A grid computing system is dedicated to support computation across a variety of administrative domains, which makes it different from the traditional distributed computing system.

III. THE DATA MINING MODEL OF CLOUD COMPUTING

Cloud security, by definition, is a new term that evolved from a "cloud computing". "Cloud Security Cloud Security" through the mesh of a large number of customer-side software on the network behaviour anomaly detection, access to the Internet, Trojans, malicious programs, pushed to the Server side for automatic analysis and processing, and then viruses and Trojan solutions distributed to each client [2]. Cloud computing is Internet-based computing, whereby shared resources, software and information are provided to computers and other devices on-demand, like the electricity grid. Cloud computing (cloud computing), is an Internet-based method of calculation, this way, the shared hardware and software resources and information on demand for computers and other devices. The cloud is in fact the network, a metaphor of the Internet. The core idea of cloud computing, a large number of unified management and scheduling for network-connected computing resources to form a pool of computing resources on demand service to users. Data mining is more complex than information search. Massive data processing to achieve high performance machine or large-scale computing devices, cloud computing-based data mining to better achieve their goals. The cloud computing model has many advantages, low cost, fault tolerance, the calculation speed; convenient program development nodes increase more easily. Can be said that cloud computing is generally applicable to data mining is more ideal computing model, from huge amounts of data to find useful and understandable knowledge of the technical means. Cloud computing is an Internet-based computing model, public participation, the computing resources (including computing power, storage capacity and interactive capabilities, etc.) is a dynamic, scalable, virtualization, and the way services. : Dynamic cloud computing and scalable computing power and efficient mass data mining has brought the possibility of; cloud computing environment the wisdom of public participation in environment swarm intelligence groups for the study set a new data mining method; cloud calculate the service features of the mass of data mining possible. Data mining in these parts have a wide range of applications, such as the web store page to research process in the page ranking and the front-end interactive query suggestions, each of which needs the support of data mining technology. Through cloud computing, mass data storage and distribution of computing, massive data mining environment for cloud computing provides new ways and means to effectively solve the distributed storage of massive data mining and efficient computing [4]. Carry out data mining method based on cloud computing, more and more complex mass data mining can provide new theoretical and support tools. Extension of cloud computing will drive the Internet and technological achievements in the public service is to promote the depth of information resources sharing and sustainable use of new methods and new ways of traditional data mining.

IV. THE MAPREDUCE PROGRAMMING FRAMEWORK

On top of the DFS, many different higher-level programming frameworks have been developed. The most commonly implemented programming framework is the MapReduce framework [4,11,12]. MapReduce is an emerging programming framework for data-intensive applications proposed by Google. MapReduce borrows ideas from functional programming [12], where the programmer defines Map and Reduce tasks to process large sets of distributed data. Implementations of MapReduce [11] enable many of the most common calculations on large-scale data to be performed on computing clusters efficiently and in a way that is tolerant of hardware failures during computation. However MapReduce is not suitable for online transactions [11,12].

The key strengths of the MapReduce programming framework are the high degree of parallelism combined with the simplicity of the programming framework and its applicability to a large variety of application domains [4,11]. This requires dividing the workload across a large number of machines. The degree of parallelism depends on the input data size. The map function processes the input pairs (key1, value1) returning some other intermediary pairs (key2, value2). Then the intermediary pairs are grouped together according to their key. The reduce function will output some new key-value pairs of the form (key3, value3). Figure shows an example of a MapReduce algorithm used to count words in a file. In this example the map input key is the provided data chunk with a value of 1. The map output key is the word itself and the value is 1 every time the word exists in the processed data chunk. The reducers perform the aggregation of the key-values pair output from the maps and output a single value for every key, which in this case is a count for every word. High performance is achieved by breaking the processing into small units of work that can be run in parallel across potentially hundreds or thousands of nodes in the cluster. Programs written in this functional style are automatically parallelized and executed on a large cluster of commodity machines. This allows programmers without any experience with parallel and distributed systems to easily utilize the resources of a large distributed system [3,4]. MapReduce programs are usually written in Java; however they can also be coded in languages such as C++, PHP Perl, Python, Ruby, R, etc. These programs may process data stored in different file and database systems.

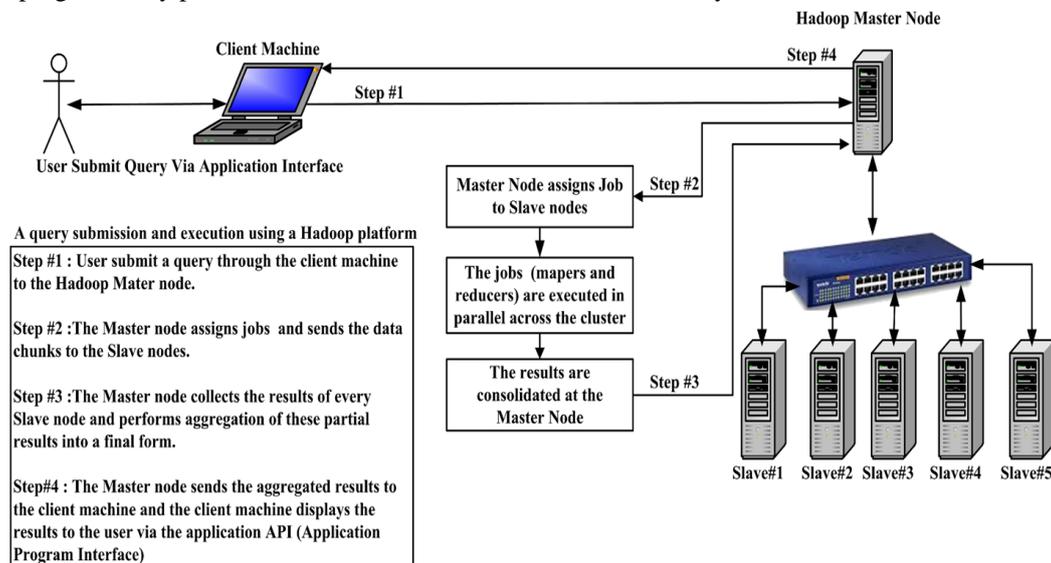


Fig.2 The Architecture of the Hadoop HDFS cluster

The hadoop platform Hadoop [13-15] is an open source software implementation of the MapReduce framework for running applications on large clusters built of commodity hardware from Apache [16]. Hadoop is a platform that provides both distributed storage and computational capabilities. Hadoop was first comprehended to fix a scalability issue that existed in Nutch [15,17], an open source crawler and search engine that utilizes the MapReduce and big-table [17] methods developed by Google. Its storage and computational capabilities scale with the addition of computing nodes to a Hadoop cluster, and can reach volume sizes in the petabytes on clusters with thousands of nodes. Hadoop differs from other distributed system schemes in its philosophy toward data.

A traditional distributed system requires repeat transmissions of data between clients and servers [3]. This works fine for computationally intensive work, but for data-intensive processing, the size of data becomes too large to be moved around easily. Hadoop focuses on moving code to data instead of vice versa [13,14]. The client (NameNode) sends only the MapReduce programs to be executed, and these programs are usually small (often in kilobytes). More importantly, the move-code-to-data philosophy applies within the Hadoop cluster itself. Data is broken up and distributed across the cluster, and as much as possible, computation on a chunk of data takes place on the same machine where that chunk of data resides.

V. PROPOSED WORK

Data mining is a process of extracting information from the raw data and Cloud computing provides scalable and flexible infrastructure which provides everything as a service. By integrating data mining and cloud computing provides agility and quick access to technology. The traditional database system and RDBMS are not able to mine large data set and the

existing data mining algorithms are not capable to mine big data.

Hadoop is a distributed master-slave architecture that consists of the Hadoop Distributed File System (HDFS) for storage and the MapReduce programming framework for computational capabilities. The HDFS stores data on the computing nodes providing a very high aggregate bandwidth across the cluster. Traits inherent to Hadoop are data partitioning and parallel computation of large datasets. MapReduce framework [4,11,12]. MapReduce is an emerging programming framework for data-intensive applications proposed by Google. MapReduce borrows ideas from functional programming [12], where the programmer defines Map and Reduce tasks to process large sets of distributed data. Implementations of MapReduce [11] enable many of the most common calculations on large-scale data to be performed on computing clusters efficiently

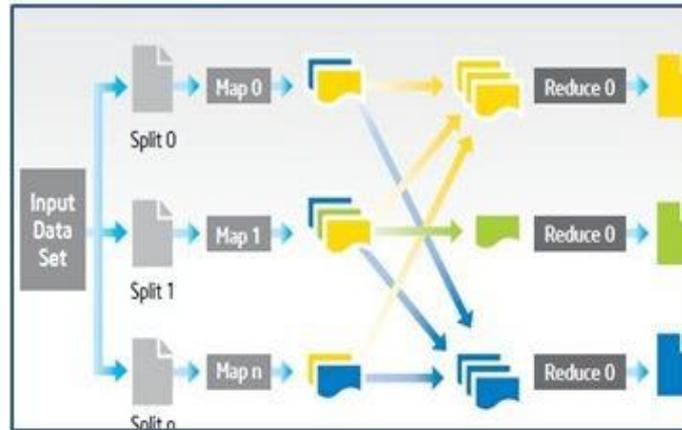


Fig.3 MapReduce Feramework

Map Reduce framework allows users to define two functions, map and reduce, to process a large number data entries. Here a scalable parallel clustering algorithm is used to overcome the problem in clustering large dataset with high dimension. The clustering is the Fuzzy C-Means (FCM) clustering algorithm which is applied to the each randomly divided set of input data[22]. Then finally the resultant cluster is obtained at the output. The fig. shown below represents the architecture of the proposed scalable parallel clustering algorithm.

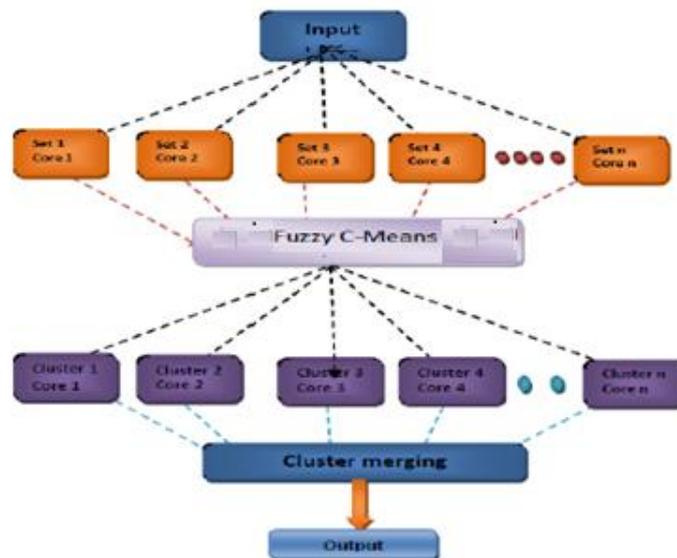


Fig.4 Parallel architecture of the proposed algorithm

VI. PARTITIONING THE INPUT LARGE DATASET

Let the input be the large dataset with a size of $M * N$. In this processing, input large dataset using fuzzy c-means clustering algorithm is difficult. So dividing the input dataset randomly in to small subsets of data with equal size will make system better [28]. So further in this proposed system the input large data set is divided into N number of subset, based on number of cores available in the system, $S=\{S_1, S_2, \dots, S_n\}$ where N is the total number of sets with equal size. Here the each subset of data is clustered into clusters using a standard and efficient clustering algorithm called Fuzzy C-Means. Programmatically used in fork method in Java. The each single data subset S consist of a vector of d measurements, where $X = (x_1, x_2, \dots)$. The attribute of an individual data set is represented as $i \times$ and d represents the dimensionality of the vector. The Fuzzy C-Means[30] is applied to the each subset of dataset for clustering the input dataset $n \times d$ in to k-clusters. Fuzzy C-Means clustering method is applied to divided subset of data. The PFCM is one of the most efficient parallel clustering methods. Let the unlabelled data set is $S=\{S_1, S_2, \dots, S_n\}$, which is further clustered in to a group of k-clusters using CMeans clustering method.

VII. INTEGRATED DATA MINING AND CLOUD COMPUTING

The integrated approach of data mining and cloud computing and mining is the process of extracting structured information from unstructured or semi-structured web data sources. The application of this technology should enable that with just a few clicks one can collect the information about the end user of the application entirely. It provides technology that can handle large amount of data which cannot be processed efficiently at reasonable cost using standard technologies and techniques. It also allows the users to retrieve meaningful information from virtually integrated data warehouse that reduces the cost of infrastructure and storage. Extension of cloud computing will drive the internet and technological achievements in the public service is to promote the depth of information resources sharing and sustainable use of new methods and new ways of traditional data mining.

VIII. ADVANTAGES OF INTEGRATED DATA MINING AND CLOUD COMPUTING

The following are the advantages of the integrated data mining and cloud computing environment.

- i. Virtual m/c that can be started with short notice
- ii. Redundant robust storage
- iii. No query structured data
- iv. Message queue for communication
- v. The customer only pays for the data mining tools that he needs
- vi. The customer doesn't have to maintain a hardware infrastructure as he can apply data mining through a browser

IX. IMPLEMENTAION OF BIG DATA MINING ON HDFS, MAPREDUCE THOUGH C-MEANS

After configuration HDFS, we are configuring our application for big data mining. In this step, we are using configured an iso image file of our virtual machine. This setup can be run over any virtual platform; we need not to configure it again for other systems. Over this architecture we configure our Big Data Mining application. This application will process map reduce's huge data according to the filter criteria given to the application. We need not to use upper layer tools of HDFS architecture for further implementation. We are using lamp server which is configured on this architecture for our application for data retrieval.

Configuring Hadoop

Configuring and getting Hadoop up and running is quite straightforward. However, since this process requires editing multiple configuration and setup files, make sure that each step is properly followed.

1. Install Java

Hadoop requires Java to be installed, so let's begin by installing Java:

```
apt-get update  
apt-get install default-jdk
```

These commands will update the package information on your VPS and then install Java. After executing these commands, execute the following command to verify that Java has been installed:

2. Configure HDFS

First let's fetch Hadoop from one of the mirrors using the following command:

```
wget http://www.motorlogy.com/apache/hadoop/common/current/hadoop-2.3.0.tar.gz  
After downloading the Hadoop package, execute the following command to extract it:  
tar xzf hadoop-2.3.0.tar.gz
```

This command will extract all the files in this package in a directory named hadoop-2.3.0. For this tutorial, the Hadoop installation will be moved to the /usr/local/hadoop directory using the following command:

```
mv hadoop-2.3.0 /usr/local/hadoop
```

After completing all the configuration outlined in the above steps, the Hadoop filesystem needs to be formatted so that it can start being used. This is done by executing the following command:

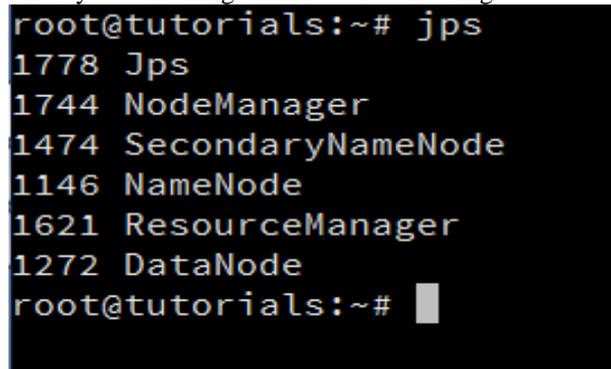
```
hdfs namenode -format
```

3. Start Hadoop

All that remains to be done is starting the newly installed single node cluster:

```
start-dfs.sh
```

Executing this command should show you something similar to the following:



```
root@tutorials:~# jps  
1778 Jps  
1744 NodeManager  
1474 SecondaryNameNode  
1146 NameNode  
1621 ResourceManager  
1272 DataNode  
root@tutorials:~# █
```

Fig 5 Showing Name and Data Node of Hadoop

After configuration hadoop, we are configuring our application for big Data Analysis. In this step, we are using configured an iso image file of our virtual machine. This setup can be run over any virtual platform; we need not to configure it again for other systems.

Starting the multi-node cluster

Starting the cluster is performed in 2 steps.

- i) Starting the HDFS Daemons – the NameNode daemon is started on the master (Hadoopmaster) and the DataNode daemons are started on all slaves (in our case Hadoopmaster and hadoopnode).
- ii) Start the MapReduce Daemons - The JobTracker is started on the master (Hadoopmaster), and the TaskTracker daemons are started on all slaves (in our case Hadoopmaster and hadoopnode)

HDFS Daemon.

Start hdfs by executing the following command on the master (Hadoopmaster). This will bring up the HDFS with the NameNode up on the master and, the DataNode on the machine listed in the conf/slaves file.

bin/start-dfs.sh

```
hduser@Hadoopmaster:~$ start-dfs.sh
Warning: $HADOOP_HOME is deprecated.

starting namenode, logging to /usr/local/hadoop/libexec/./logs/hadoop-hduser-na
menode-Hadoopmaster.out
hadoopnode: starting datanode, logging to /usr/local/hadoop/libexec/./logs/hado
op-hduser-datanode-hadoopnode.out
Hadoopmaster: starting datanode, logging to /usr/local/hadoop/libexec/./logs/ha
dooop-hduser-datanode-Hadoopmaster.out
Hadoopmaster: starting secondarynamenode, logging to /usr/local/hadoop/libexec/.
./logs/hadoop-hduser-secondarynamenode-Hadoopmaster.out
```

Fig 5 Showing HDFS Daemon

Mapred Daemon.

Start mapreduce by executing the following command on the master (Hadoopmaster). This will bring up the MapReduce cluster with the JobTracker up on the master and, the TaskTracker on the machine listed in the conf/slaves file. bin/start-mapred.sh

```
hduser@Hadoopmaster:~$ start-mapred.sh
Warning: $HADOOP_HOME is deprecated.

starting jobtracker, logging to /usr/local/hadoop/libexec/./logs/hadoop-hduser-
jobtracker-Hadoopmaster.out
Hadoopmaster: starting tasktracker, logging to /usr/local/hadoop/libexec/./logs
/hadoop-hduser-tasktracker-Hadoopmaster.out
hadoopnode: starting tasktracker, logging to /usr/local/hadoop/libexec/./logs/h
adoop-hduser-tasktracker-hadoopnode.out
hduser@Hadoopmaster:~$
```

Fig 6 Showing MapReduce Daemon

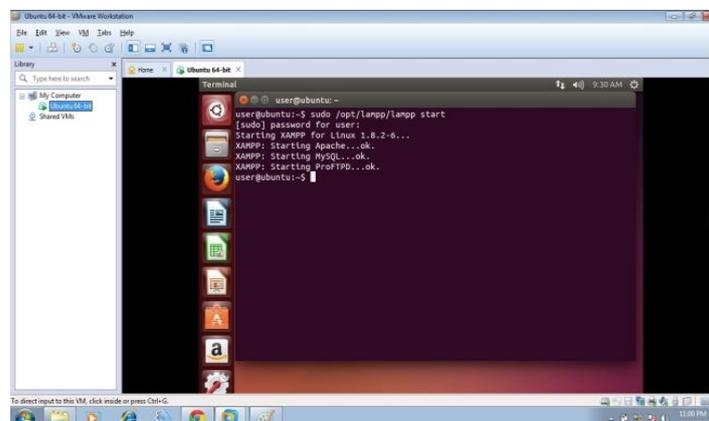


Fig.7 Configured LAMPP and Services

To start and run the mining application, we use browser. Call the application start page by giving the url on address bar. After successfully call and running the index page, the application visualize as the following screen.

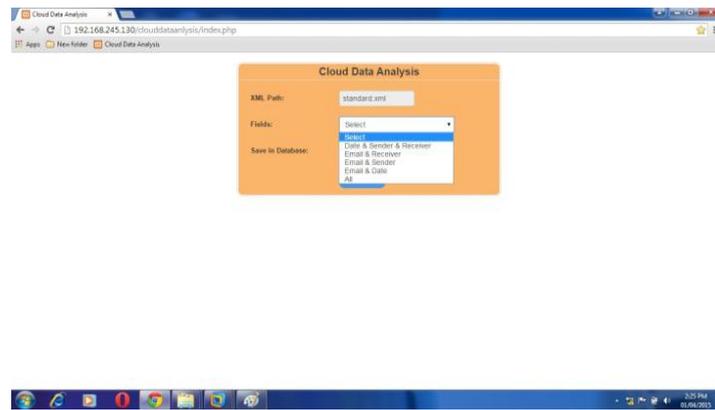


Fig.8 Big Data Mining Application

X. CONCLUSION

This paper presents a review of need data mining services in cloud computing along with a case study on the integrated approach of data mining and cloud computing and mining. The data mining in cloud allows organization to centralize the management of software and data storage with assurance of efficient, reliable and secure services for their users. The implementation of data mining techniques through cloud computing will allow the users to retrieve meaningful information from virtually integrated data warehouse that reduces the costs of infrastructure and storage. This approach also reduces the barriers that keep small companies from benefiting of the data mining instruments. The emergence of cloud computing brings new ideas for data mining. It increases the scale of processing data.

This approach is deployed in HDFS, MapReduce and C-Means big data mining algorithm framework of Hadoop. Big data discloses the limitations of existing data mining techniques, resulted in a series of new challenges related to big data mining. Big data mining is a promising research area, still in its infancy. In spite of the limited work done on big data mining so far, we believe that much work is required to overcome its challenges related to heterogeneity, scalability and execution speed.

REFERENCES

- [1] Shuman S: Structure, mechanism, and evolution of the mRNA capping apparatus. *Prog Nucleic Acid Res Mol Biol* 2000.
- [2] Rajaraman A, Ullman JD: *Mining of Massive Datasets*. Cambridge – United Kingdom: Cambridge University Press; 2012.
- [3] Coulouris GF, Dollimore J, Kindberg T: *Distributed Systems: Concepts and Design*: Pearson Education; 2005.
- [4] de Oliveira Branco M: *Distributed Data Management for Large Scale Applications*. Southampton – United Kingdom: University of Southampton; 2009.
- [5] Raghupathi W, Raghupathi V: Big data analytics in healthcare: promise and potential. *Health Inform Sci Syst* 2014, 2(1):3.
- [6] Bell DE, Raiffa H, Tversky A: *Descriptive, normative, and prescriptive interactions in decision making*. *Decis Mak* 1988
- [7] Foster I, Kesselman C: *The Grid 2: Blueprint for a new Computing Infrastructure*. Houston – USA: Elsevier; 2003.
- [8] Owens JD, Houston M, Luebke D, Green S, Stone JE, Phillips JC: GPU computing. *Proc IEEE* 2008, 96(5):879–899.
- [9] Satish N, Harris M, Garland M: Designing efficient sorting algorithms for manycore GPUs. In *Parallel & Distributed Processing, 2009 IPDPS 2009 IEEE International Symposium on: 2009, IEEE; 2009:1–10*.
- [10] He B, Fang W, Luo Q, Govindaraju NK, Wang T: Mars: a MapReduce framework on graphics processors. In *Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques: 2008*,
- [11] Dean J, Ghemawat S: MapReduce: simplified data processing on large clusters. *Commun ACM* 2008, 51(1):107–113.
- [12] Peyton Jones SL: *The Implementation of Functional Programming Languages (Prentice-Hall International Series in Computer Science)*. New Jersey – USA: Prentice-Hall, Inc; 1987.
- [13] Bryant RE: *Data-intensive supercomputing: The case for DISC*. Pittsburgh, PA – USA: School of Computer Science, Carnegie Mellon University; 2007:1–20.
- [14] White T: *Hadoop: The Definitive Guide*. Sebastopol – USA: “O’Reilly Media, Inc.”; 2012.
- [15] Shvachko K, Kuang H, Radia S, Chansler R: The hadoop distributed file system. In *Mass Storage Systems and Technologies (MSST), 2010 IEEE 26th Symposium on: 2010, IEEE; 2010:1–10*.
- [16] The Apache Software Foundation. [<http://apache.org/>]
- [17] Olson M: Hadoop: Scalable, flexible data storage and analysis. *IQT Quart* 2010, 1(3):14–18.
- [18] Xiaojing J: Google Cloud Computing Platform Technology Architecture and the Impact of Its Cost. In *2010 Second WRI World Congress on Software Engineering: 2010; 2010:17–20*.

- [19] Thusoo A, Sarma JS, Jain N, Shao Z, Chakka P, Anthony S, Liu H, Wyckoff P, Murthy R: Hive: a warehousing solution over a map-reduce framework. Proc VLDB Endowment 2009, 2(2):1626–1629.
- [20] Olston C, Reed B, Srivastava U, Kumar R, Tomkins A: Pig latin: a not-so-foreign language for data processing. In Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data: 2008, ACM; 2008:1099–1110.
- [21] S.Prabha and P.Kola Sujatha, “REDUCTION OF BIG DATA SETS USING FUZZY CLUSTERING” International Journal of Advanced Research in Computer Engineering & Technology (IJARCET) Volume 3 Issue 6, June 2014.
- [22] R Madhuri, M R Murty, JVR Murthy, PVGD Prasad Reddy and S.C Satapathy “Cluster analysis on different data sets using k-modes and k-prototype algorithms,” ICT and Critical Infrastructure: Proceedings of the 8th Annual Convention of Computer Society of India, Vol. II, pp 137-144. Springer, 2014.
- [23] X.F Jiang “Application of parallel annealing particle clustering algorithm in data mining,” TELKOMNIKA Indonesian Journal of Electrical Engineering, Vol. 12, No. 3, pp. 2118-2126, 2014.
- [24] R. Krishnapuram and J. M. Keller, —A possibilistic approach to clustering, IEEE Transactions on Fuzzy Systems, vol. 1, p.10.1109/91.227387, 1993.
- [25] Janardhan. N, T. Sree Pravallika, Sowjanya Gorantla, “An efficient approach for integrating data mining into cloud computing”, International Journal of Computer Trends and Technology (IJCTT) - volume4 Issue5–May 2013
- [26] Daniel J. Abadi, Yale University, DataManagement in the Cloud: Limitations and Opportunities, Bulletin of the IEEE Computer Society Technical Committee on Data Engineering 2009.
- [27] Naskar Ankita, Mrs. Mishra Monika R., “Using Cloud Computing To Provide Data Mining Services”, International Journal Of Engineering And Computer Science ISSN:2319-7242 Volume 2 Issue 3 March 2013 Page No. 545-550.
- [28] Zeba Qureshi, Jaya Bansal, Sanjay Bansal, A Survey on Association Rule Mining in Cloud Computing, paper of ijetaeISSN 2250-2459, Volume 3, Issue 4, April 2013).
- [29] Robert Vrbić, “Data Mining and Cloud Computing”, Journal of Information Technology and Applications, Volume 2 December 2012.
- [30] B R Prakash, Dr. M. Hanumanthappa, Issues and Challenges in the Era of Big Data Mining, IJETTCS- ISSN 2278-6856, Volume 3, Issue 4 July-August 2014