Hadoop Cluster with FPGA-based Hardware Accelerators for K-means Clustering Algorithm

Ching-Che Chung, *Member, IEEE*, and Yu-Hsin Wang
Department of CSIE, National Chung Cheng University, Taiwan.

***Abstract*-- In this paper, the implementation of the K-means clustering algorithm on a Hadoop cluster with FPGA-based hardware accelerators is presented. The proposed design follows MapReduce programming model and uses Hadoop distribution file system (HDFS) for storing large dataset. The proposed FPGA-based hardware accelerator for speed up the K-means clustering algorithm is implemented on Xilinx VC707 evaluation boards (EVBs). There are four computers in the proposed Hadoop cluster, one computer is Master Node, and the other three computers are Slave Nodes. The Slave Nodes communicate with VC707 EVBs through Gigabit Ethernet. The experimental results show that for clustering 125 million three-dimensional input dataset, the proposed design can achieve 4x speedup than the Hadoop cluster without FPGA-based hardware accelerators.**

I. INTRODUCTION

Nowadays, big data analytics has become the focus of research in the various fields of government, manufacturing, healthcare, media, and science because of the growing popularity of internet of things (IoT). In fact, all of the industries have to confront the issues of big data analytics. The properties of big data include variety, volume, velocity and value. The "4Vs" is widely applied to the definition of big data [1]. Due to these properties, big data is not easy to be analyzed using a single computer. Many software framework developments with high scalability and fault tolerance, for example, Hadoop and Spark, are provided to handle massive data.

Hadoop [3] is a software framework that enables massive data storage and distributed processing over large number of computing servers. Hadoop distributed file system (HDFS) is a distributed file system which provides high throughput access to the application data. In HDFS, a file will be split into one or more blocks, and each block will have several replications on different DataNode to prevent missing data. MapReduce is a software framework for easily writing applications which process vast amounts of data in parallel on large computing clusters in a reliable and fault-tolerant manner.

Recently, ARM-based Hadoop clusters with FPGA-based hardware accelerators to improve the performance in big data analytics are proposed. In ZCluster [2], Zynq platform which integrates the ARM processor and the FPGA in the single chip is used as the Slave Node. Then, the FPGA can act as a co-processor to share the loading of the ARM processor. Thus, the execution time of applications can be reduced. However, in addition to the jobs executed in the FPGA, some irregular jobs of the applications still need to run with ARM processor. Thus, if only a small portion of jobs can be executed by the FPGA, the relatively slow performance of the ARM processor as compared to the Intel I7 processors will be the problem.

In this paper, we use Xilinx VC707 evaluation boards (EVBs) [5] to communicate with the host computers through Gigabit Ethernet. The host computer with high performance processors can share the loading of applications by sending the jobs to the EVBs. We use K-means clustering algorithm as an example to demonstrate the speed up of the Hadoop cluster with FPGA-based hardware accelerators.

II. PROPOSED ARCHITECTURE

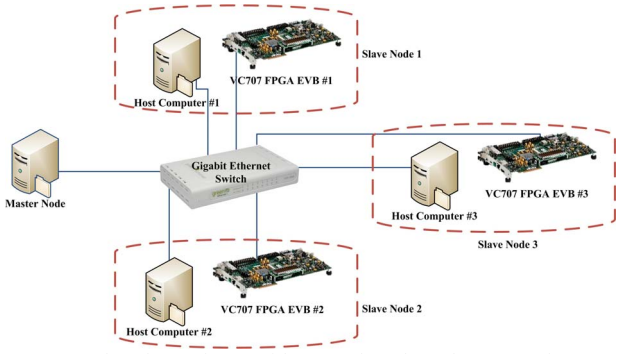

Fig. 1 Proposed Hadoop Cluster with FPGA-based Hardware Accelerators.

The proposed Hadoop cluster with FPGA-based hardware accelerators is shown in Fig. 1. In the host computers, Hadoop [3] with version: 2.7.1 was installed. The specifications of the host computer are CPU: Intel(R) Core(TM) i7-4770 @ 3.4GHz with a solid-state disk (SSD) and the operating system is 64-bit CentOS 6.7. The Hadoop cluster consists of one Master Node and three Slave Nodes, and in each Slave Node, it consists of one host computer and one VC707 EVB. In the Slave Node, the host computer communicates with VC707 EVB through Gigabit Ethernet.

The system behavior is described as following. Firstly, we write the input dataset to HDFS and execute the MapReduce program in the Master Node. Then, in map stage, Hadoop sends map tasks to the appropriate servers in the Hadoop cluster. In each map task, the host computer reads the dataset from HDFS, and it wraps nodes data, cluster centers and control signals into transmission packets. Subsequently, these packets are sent to the VC707 EVBs. The VC707 EVB receives packets and the hardware accelerator implemented in the FPGA will perform K-means clustering algorithm with parallel hardware circuits to group nodes into different clusters. After that, the hardware accelerator implemented in the FPGA calculates the number of nodes in each cluster and

This work was supported in part by the Ministry of Science and Technology of Taiwan, under Grant MOST103-2221-E-194-063-MY3.

returns the partial sum of each dimension in the clusters to the host computer. Then, we generate the index of the cluster as key and the partial sum of each dimension in each cluster and the number of nodes in each cluster as value. Finally, in reduce stage, Hadoop sends reduce tasks to the appropriate servers in the Hadoop cluster. In reduce task, according to the <key, value> from map stage, the new cluster center can be calculated for the next iteration.

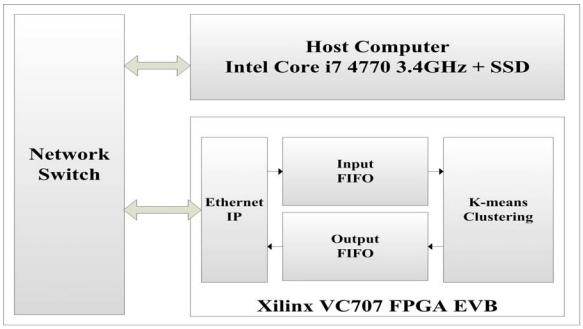

Fig. 2 Block diagram of the Slave Node.

Fig. 2 shows the block diagram of the Slave Node, the Slave Node is composed of one host computer and one VC707 EVB. The modules implemented in the FPGA including Ethernet physical IP, input FIFO, output FIFO, mixed-mode clock manager (MMCM), and K-means clustering circuits. The host computer wraps data into packets and sends them to VC707 EVB through a Gigabit Ethernet switch. Then, the VC707 EVB sends back the computation results to the host computer. Each packet has the maximum size of 1500 bytes. Therefore, 115 nodes three-dimensional single precision floating point coordinate values can be sent to the VC707 EVB in one packet. The K-means clustering circuit can handle these nodes simultaneously and returns the computation results. Finally, the host computer can calculate new centers for the next iteration.

In K-means clustering circuit, the inputs are 115 nodes three-dimensional single precision floating point coordinate values and the initial center coordinate values for four clusters from the host computer. When the K-means clustering circuit receives the coordinate values of 115 nodes, it calculates the Euclidean distances of each node to the cluster centers with parallel hardware circuits and finds the shortest distances to the cluster centers. After that, the nodes can be grouped into four clusters. When the input nodes are all grouped into different clusters, the number of nodes in each cluster and the partial sum of the coordinate values of each dimension in each cluster are sent back to the host computer. Then, the host computer sends another 115 nodes to the VC707 EVB until all input dataset are grouped by the proposed FPGA-based hardware accelerator. In addition, the host computer calculates new centers for the next iteration by accumulating the partial sum and the number of nodes in each cluster which are sent by the VC707 EVB.

III. Experimental Results

In the proposed Hadoop cluster with FPGA-based hardware accelerators, the transmission data rate from the host computer to the VC707 EVB and from the VC707 EVB to the host computer are 436 Mbps and 125 Mbps, respectively. Fig. 3 shows the total time required for K-means clustering algorithm with different size of input dataset. The proposed Hadoop streaming with FPGA-based hardware accelerators can achieve 4x speedup than the Hadoop cluster using Mahout machine learning libraries [4]. Fig. 4 shows the total time required for K-means clustering algorithm with different number of Slave Nodes. When the number of Slave Nodes is increased, the total time required for K-means clustering algorithm can be reduced accordingly. Moreover, it also shows the performance saturation when the number of Slave Nodes is increased from two to three.

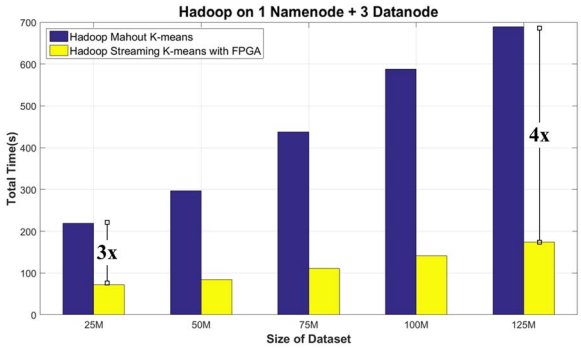

Fig. 3 Hadoop cluster with FPGA-based accelerators.

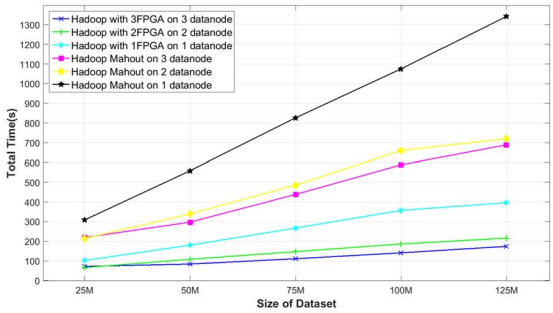

Fig. 4 Hadoop cluster with 1/2/3 Slave Nodes.

IV. Conclusion

In this paper, the proposed Hadoop cluster with FPGA-based hardware accelerators for K-means clustering algorithm can share the loading in the Slave Nodes to the VC707 EVBs. As a result, the proposed high-scalable heterogeneous computing solution is suitable to be applied to different machine learning algorithms for big data analytics.

References

[1] Han Hu, et al., "Toward scalable systems for big data analytics: a technology tutorial," *IEEE Access*, vol. 2, pp. 652-687, Jul. 2014.

[2] Zhongduo Lin and Paul Chow, "ZCluster: a Zynq-based Hadoop cluster," *in Proc. International Conference on Field-Programmable Technology (FPT)*, Dec. 2013, pp. 450-453.

[3] Apache Hadoop, available: https://hadoop.apache.org/.

[4] Apache Mahout, available: https://mahout.apache.org/.

[5] VC707 Evaluation Board for the Virtex-7 FPGA User Guide, Xilinx Inc., available: http://www.xilinx.com/support/documentation/boards_and_kits/vc707/ug885_VC707_Eval_Bd.pdf.