Research Repository UCD

Title	Performance Evaluation of a Distributed Clustering Approach for Spatial Datasets
Authors(s)	Bendeche, Malika, Le-Khac, Nhien-An, Kechadi, Tahar
Publication date	2017-08-20
Publication information	Bendeche, Malika, Nhien-An Le-Khac, and Tahar Kechadi. "Performance Evaluation of a Distributed Clustering Approach for Spatial Datasets." Springer, 2017.
Conference details	AusDM 2017: 15th Australasian Conference, Melbourne, VIC, Australia, 19-20 August 2017
Series	Communications in Computer and Information Science book series (CCIS), Volume 845
Publisher	Springer
Item record/more information	http://hdl.handle.net/10197/10848
Publisher's statement	The final publication is available at www.springerlink.com .
Publisher's version (DOI)	10.1007/978-981-13-0292-3_3

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

Performance Evaluation of a Distributed Clustering Approach for Spatial Datasets

Malika Bendeche¹, Nhien-An LE-KHAC², and M-Tahar Kechadi¹

¹ Insight centre for Data Analytics
University College Dublin,
O'Brien building, Centre East
Belfield, Dublin 04,
`malika.bendeche@gamil.com`,
`tahar.kechadi@ucd.ie`,
² University College Dublin
Belfield, Dublin 04
`an.lekhac@ucd.ie`

Abstract. The analysis of big data requires powerful, scalable, and accurate data analytics techniques that the traditional data mining and machine learning do not have as a whole. Therefore, new data analytics frameworks are needed to deal with the big data challenges such as volumes, velocity, veracity, variety of the data. Distributed data mining constitutes a promising approach for big data sets, as they are usually produced in distributed locations, and processing them on their local sites will reduce significantly the response times, communications, etc. In this paper, we propose to study the performance of a distributed clustering, called Dynamic Distributed Clustering (DDC). DDC has the ability to remotely generate clusters and then aggregate them using an efficient aggregation algorithm. The technique is developed for spatial datasets. We evaluated the DDC using two types of communications (synchronous and asynchronous), and tested using various load distributions. The experimental results show that the approach has super-linear speed-up, scales up very well, and can take advantage of the recent programming models, such as MapReduce model, as its results are not affected by the types of communications.

Keywords: Distributed data mining, distributed computing, synchronous communication, asynchronous communication, spacial data mining, super-speedup.

1 Introduction

Nowadays big data is becoming a commonplace. It is generated by multiple sources at rapid pace, which leads to very large data volumes that need to be stored, managed, and analysed for useful insights. From organisations point of view, it is not the size of the generated data which is important. It is what we learn from it that matters, as this may help understanding the behaviour of the

system that is governed by this data or help to make some key decisions, etc. To extract meaningful value from big data, we need appropriate and efficient mining and analytics techniques to analyse it. One of the most powerful and common approaches of analysing datasets for extracting useful knowledge is clustering. Clustering has a wide range of applications and its concept is so interesting that numerous algorithms for various types of data have been proposed and implemented. However, big data come up with new challenges, such as large volumes, velocity, variety, and veracity, that the majority of popular clustering algorithms are inefficient at very large scale. This inefficiency can be that the final results are not satisfactory or the algorithm has high complexity which requires large computing power and response time to produce final results. There are two major categories of approaches to deal with computational complexity of these clustering algorithms: 1) the first category consists of reducing the size of the initial dataset. One can use either sample-based techniques or dimensionality reduction techniques. The second category consists of using parallel and distributed computing to speed up the response time. In this case we can try to parallelise or model the algorithm in the form of a client-server model using MapReduce mechanism. However, these algorithms are inherently difficult to parallelise, and designing an efficient distributed version of the algorithm is not straightforward either. This is due to the fact that the processing nodes, either in the parallel or distributed versions, need to communicate and coordinate their efforts in order to obtain the same results. These communications are extremely expensive and can cancel the benefit of the parallelised version. To deal with these challenging issues, we propose to study a distributed approach that takes advantage of parallel and distributed computing power, while getting ride of the drawbacks of the previous methods. In addition, one of the main advantages of our approach is that it can be used as a framework for all clustering algorithms. In other words, while it is well known that there is no clustering algorithm that can universally be used to cluster every dataset of a given application, our approach can be used for all algorithms or a set of algorithms to derive a distributed clustering approach for a given data having specific characteristics.

The proposed approach has two main phases: the first phase, based on the SPMD (Single Program Multiple Data) paradigm, consists of dividing the datasets into K partitions, where K is the number of processing nodes. Then, for each partition we cluster its data into C_i clusters. This phase is purely parallel, as each processing node executes a clustering algorithm on its data partition independently of the others. The obtained clusters on each node are called local clusters. This phase does not require any communications, and in addition, in the majority of applications the data is collected by various sources, which are geographically distributed. Therefore, the data is already partitioned. All required is to cluster locally the data. The second phase consists of aggregating (or merging) the local clusters to obtain global clusters by merging overlapping clusters. In order to determine whether two local clusters belonging to two different nodes are overlapping or not, one needs to exchange the local clusters between the nodes. This operation is extremely expensive when the dataset is

very large. The main idea of our approach is to minimise the data exchange while maximising the quality of the global clusters. The method used to aggregate spatial local clusters into global clusters allows only to exchange about 2% of the original datasets [29], which is highly efficient. In this paper, we want to study the performance of such distributed clustering technique by calculating its speedup compared to the sequential version of the algorithm, its scalability, its communication overheads, and its complexity in general.

The rest of the paper is organised as follows: In the next section we will give an overview of the state-of-the-art of parallel and distributed data mining techniques and discuss their limitations. Then we will present in more details the proposed distributed framework and its concepts in Section 3. In section 4, we evaluate its performance based on two types of implementations; synchronous and asynchronous communications. In section 5, we discuss the experimental results based on speedup, scalability, communication overheads, and compare the two implementation models; synchronous and asynchronous. Finally, we conclude in Section 6.

2 Related Work

Distributed Data Mining (DDM) is a line of research that has attracted much interest in recent years [26]. DDM was developed because of the need to process data that can be very large or geographically distributed across multiple sites. This has two advantages: first, a distributed system has enough processing power to analyse the data within a reasonable time frame. Second, it would be very advantageous to process data on their respective sites to avoid the transfer of large volumes of data between the site to avoid heavy communications, network bottlenecks, etc.

DDM techniques can be divided into two categories based on the targeted architectures of computing platforms [37]. The first, based on parallelism, uses traditional dedicated and parallel machines with tools for communications between processors. These machines are generally called super-computers. The second category targets a network of autonomous machines. These are called distributed systems, and are characterised by distributed resources, low-speed network connecting the system nodes, and autonomous processing nodes which can be of different architectures, but they are very abundant [24]. The main goal of this category of techniques is to distribute the work among the system nodes and try to minimise the response time of the whole application. Some of these techniques have already been developed and implemented in [1, 35].

However, the traditional DDM methods are not always effective, as they suffer from the problem of scaling. One solution to deal with large scale data is to use parallelism, but this is very expensive in terms of communications and processing power. Another solution is to reduce the size of training sets (sampling). Each system node generates a separate sample. These samples will be analysed using a single global algorithm [38, 27]. However, this technique has a disadvantage that the sampling in this case is very complex and requires

many communications between the nodes which may impact on the quality of the samples and therefore the final results. This has led to the development of techniques that rely on ensemble learning [33, 4]. These new techniques are very promising, as each technique of the ensemble network attempts to learn from the data and the best or compromised results of the network will emerge as the winner. Integrating ensemble learning methods in DDM framework will allow to deal with the scalability problem, as it is the case of the proposed approach.

Clustering algorithms can be divided into two main categories, namely partitioning and hierarchical. Different elaborated taxonomies of existing clustering algorithms are given in the literature. Many parallel clustering versions based on these algorithms have been proposed in the literature [2, 14, 18, 22, 23, 15, 36]. These algorithms are further classified into two sub-categories. The first consists of methods requiring multiple rounds of message passing. They require a significant amount of synchronisations and data exchange. The second sub-category consists of methods that build local clustering models and send them to a central site to build global models [29].

In [14] and [15], message-passing versions of the widely used K-Means algorithm were proposed. In [18] and [36], the authors dealt with the parallelisation of DBSCAN; density-based clustering algorithm. In [22] a parallel message passing version of the BIRCH algorithm was presented. A parallel version of a hierarchical clustering algorithm, called MPC for Message Passing Clustering, which is especially dedicated to Microarray data was introduced in [23]. Most of the parallel approaches need either multiple synchronisation constraints between processes or a global view of the dataset, or both [2]. All these approaches deal with the parallelisation of the sequential version of the algorithm by trying phases of the algorithm which can be executed in parallel by several processors. However, this requires many synchronisations either to access shared data (for the shared memory model) or communications (for message passing model). In some algorithm these synchronisations and communications are extremely expensive and it is not worth parallelising them. This approach is not usually scalable.

In [9] a client-server model is adopted, where the data is equally partitioned and distributed among the servers, each of which computes the clusters locally and sends back the results to the master. The master merges the partially clustered results to obtain the final results. This strategy incurs a high communication overhead between the master and slaves, and a low parallel efficiency during the merging process. Other parallelisations using a similar client-server model include [3, 11, 12, 21, 25, 39]. Among these approaches, various programming mechanisms have been used, for example, a special parallel programming environment, called skeleton based programming in [12] and parallel virtual machine in [25]. A Hadoop-based approach is presented in [21].

Another approach presented in [2] also applied a merging of local models to create the global models. Current approaches only focus on either merging local models or mining a set of local models to build global ones. If the local models cannot effectively represent local datasets then global models accuracy

will be very poor [29]. In addition, both partitioning and hierarchical categories have some issues which are very difficult to deal with in parallel versions. For the partitioning class, it needs the number of clusters to be fixed in advance, while in the majority of applications the number of classes is not known in advance. For the hierarchical clustering algorithms, they have the issue of stopping conditions for clustering decomposition, which is not an easy task and mainly in distributed versions.

3 Dynamic Distributed Clustering

Dynamic Distributed Clustering (DDC) model is introduced to deal with the limitations of the parallel and master-slave models. DDC combines the characteristics of both partitioning and hierarchical clustering methods. In addition, it does neither inherit the problem of the number of partitions to be fixed in advance nor the problem of stopping conditions. It is calculated dynamically and generates global clusters in a hierarchical way. All these features look very promising and some of them have been thoroughly studies in [7], such as the dynamic calculation of the number of the clusters and the accuracy of the final clustering, in this study one wants to show the effect of the communications on the response time, the communication model used, the scalability of the approach, and finally its performance in terms of speed up compared to the sequential version. In this paper we will focus on

- Synchronous and asynchronous communications, as this approach can be implemented either with synchronous or asynchronous communications. Both implementations produce the same results.
- The speed-up of the DDC approach using DBSCAN as the basic algorithm for clustering the partitions. This algorithm is known to have non-polynomial complexity ($O(n^2)$).
- Scalability of the approach as the size of the dataset increases.

We start by briefly explaining the algorithm and then present a performance and evaluation model for the approach.

The DDC approach has two main phases. In the first phase, we cluster the datasets located on each processing node and select good local representatives. All local clustering algorithms are executed in parallel without communications between the nodes. As DBSCAN is the basic algorithm for clustering local datasets, we can reach a super linear speed-up of p^2 , where p is the number of processing nodes. The second phase collects the local clusters from each node and affects them to some special nodes in the system; called leaders. The leaders are elected according to their characteristics such as capacity, processing power, connectivity, etc. The leaders are responsible for merging the local clusters. In the following we explain how the local clusters are represented and merged to generate global clusters.

Fig. 1. An overview of the DDC Approach.

3.1 Local Models

The local clusters are highly dependent on the clustering techniques used locally in each node. For instance, for spatial datasets, the shape of a cluster is usually dictated by the technique used to obtain them. Moreover, this is not an issue for the first phase, as the accuracy of a cluster affects only the local results of a given node. However, the second phase requires sending and receiving all local clusters to the leaders. As the whole data is very large, this operation will saturate very quickly the network. So, we must avoid sending all the original data through the network. The key idea of the DDC approach is to send only the cluster's representatives, which constitute between 1% and 2% of the whole data. The cluster representatives consist of the internal data representatives plus the boundary points of the cluster.

There are many existing data reduction techniques in the literature. Many of them are focusing only on the dataset size. For instance, they try to reduce the storage capacity without paying attention to the knowledge contained in the data. In [30], an efficient reduction technique has been proposed; it is based on density-based clustering. Each cluster is represented by a set of carefully selected data-points, called representatives. However, selecting representatives is still a challenge in terms of quality and size [28, 29].

The best way to represent a spatial cluster is by its shape and density. The shape of a cluster is represented by its boundary points (called contour) (see Figure 1). Many algorithms for extracting the boundaries from a cluster can be found in the literature [19, 10, 31, 17, 32]. We use an algorithm based on triangulation to generate the clusters' boundaries [16]. It is an efficient algorithm for constructing non-convex boundaries. It is able to accurately characterise

the shape of a wide range of different point distributions and densities with a reasonable complexity of $\mathcal{O}(n \log n)$.

3.2 Global Models

The global clusters are generated in the second phase of the DDC. This phase is also executed in a distributed fashion but, unlike the first phase, it has communications overheads. This phase consists of two main steps, which can be repeated until all the global clusters were generated. First, each leader collects the local clusters of its neighbours. Second, the leaders will merge the local clusters using the overlay technique. The process of merging clusters will continue until we reach the root node. The root node will contain the global clusters (see Figure 1).

In DDC we only exchange the boundaries of the clusters. The communications can be synchronous or asynchronous. We implemented this phase using both types of communications. An evaluation model is presented in the next Section.

The pseudo codes of the two phases of the DDC framework are given in Algorithm 1 and Algorithm 2.

Algorithm: DDC algorithm : Phase 1: Local clustering

Initialisation

$Node_i \in N$, N : The total nodes in the system.

input : X_i : Dataset Fragment, $Params_i$: Input parameters for the local clustering: $Params_i = (Eps_i, MinPts_i)$ for DBSCAN for example

output: C_i : Cluster's contours of $Node_i$

foreach $Node_i$ **do**

$L_i = \text{Local_Clustering}(X_i, Params_i);$
// Node_i executes a clustering algorithm locally.
 $C_i = \text{Contour}(L_i);$
// Node_i executes a contour algorithm locally.

return (C_i);

4 DDC Evaluation

In order to evaluate the performance of the DDC approach, we use different local clustering algorithms. For instance, with both K-Means [6] and DBSCAN [7,8], the DDC approach outperforms existing algorithms in both quality of its results and response time including K-Means and DBSCAN applied to the whole dataset. In this section we evaluate its speed-up, scalability, and which architecture is more appropriate to implement it. In addition, we compare DDC with the sequential version of the basic clustering algorithm used within DDC.

Algorithm: DDC algorithm : Phase 2: Merging

input : D : Tree degree, C_i : Local cluster's contours generated by $Node_i$ in the phase 1

output: $C_{G_k, level}$: Global Cluster's contours (global results, level=0)

```

repeat
    level = treeHeight;
    Nodei joins a group  $G_{K, Level}$  of  $D$  elements;
    // Nodei joins its neighbourhood
    Nodej = ElectLeaderNode( $G_{K, Level}$ );
    // Nodej is the leader of the group  $G$ 
    // In parallel
    foreach Nodei ∈  $G_{K, Level}$  do
        if ( $i <> j$ ) then
            Send ( $C_i$ , Nodej);
            // Each node sends its contours to others nodes in the
            // same group of neighbourhood
        else
            Recv ( $C \equiv (\{C_i\}, Node_i)$ );
            // If the node is the leader, it will receive the others
            // node's contours in the same group of neighbourhood
             $G_{K, Level}$  = Merge ( $C_i$ ,  $C$ );
            // Merge the overlapping contours
    level -- ;
until (level == 0);
return ( $C_{G_k, 0}$ )

```

The proposed approach is more developed for distributed systems than pure parallel systems. Therefore, it is worth analysing the benefits of using synchronous or asynchronous processing mechanism, as distributed systems are asynchronous and the blocking operations have a strong impact in communication time [34].

In synchronous model, as illustrated in Figure 2, although machines M_3 and M_4 have finished their computations before M_1 and M_2 , they can not send their results until M_1 and M_2 finish as well. In this model, Not only the computations and communications are not overlapped but also the machines which finished early wasted sometime waiting for the other to finish [34].

In asynchronous model the machines which finished early can advance to the next step. The machines manage their communications and the 1st and 2nd phase overlap. This model is much more suitable for distributed computing, where the nodes are heterogeneous and the communications are usually slow. As can be seen in the example given in Figure 3, M_3 and M_4 start merging their results before M_1 and M_2 finish their computations.

Fig. 2. Synchronous communications.

Fig. 3. Asynchronous communications.

4.1 DDC Computational Complexity

Let M be the number of nodes and n_i the dataset given to each node v_i in the system. The complexity of our approach is the sum of its components' complexity; local mining, local reduction, and global aggregation.

Phase1 - Local clustering: Let $\Gamma(n_i)$ denote the local clustering algorithm running on node (v_i) , and $\Delta(c_i)$ be the time required to execute the reduction algorithm. The cost of this phase is given by:

$$T_{Phase_1} = \max_{i=1}^M (\Gamma(n_i) + \Delta(c_i)) \quad (1)$$

Where c_i is the cluster points generated by node v_i . Note that the reduction algorithm is of complexity $\mathcal{O}(c_i \log c_i)$.

Phase2 - Aggregation: The aggregation depends on the hierarchical combination of contours of local clusters. As the combination is based on the intersection of edges from the contours, the complexity of this phase is $\mathcal{O}(w_i \log w_i + p)$. Where w_i is the total vertices of the contours by node v_i and p is the intersection points between edges of different contours (polygons).

Total complexity: The total complexity of the DDC approach, assuming that the local clustering algorithm is DBSCAN which is of complexity $\mathcal{O}(n^2)$, is:

$$T_{Total} = \mathcal{O}(n_i^2) + \mathcal{O}(c_i \log c_i) + \mathcal{O}(w_i \log w_i + p) \simeq \mathcal{O}(n_i^2) \quad (2)$$

4.2 DDC Speedup

The DDC speedup is calculated against the sequential version of the approach. The sequential version consists clustering all the data on one machine. Therefore, it does require neither reduction nor aggregation. Let T_1 be the execution time of the sequential version and T_p the execution time of the DDC on p nodes. The speedup α is given by

$$\alpha = \frac{T_1}{T_p} \quad (3)$$

Note that if the complexity of the clustering algorithm is polynomial then the optimal speedup that can be reached is P , under the condition that there is no overhead due to communications and extra work. If the complexity of the clustering algorithm is $\mathcal{O}(n^2)$ then the optimal speedup can be P^2 ; this is called *super speedup*. In the following section we will evaluate the speedup in the case of DBSCAN.

5 Experimental Results

We have implemented our approach on a distributed computing system. The distributed computing system consists of heterogeneous desktops (different CPUs, OSs, memory sizes, loads, etc.). We use JADE (Java Agent DEvelopment), as a development platform to implement the approach. JADE is based on a 2P2 communication architecture, It allows to use heterogeneous processing nodes, it is scalable, and dynamic [13, 5].

The system nodes (desktops) are connected to local area networks. This allows us to add as many nodes as required, depending on the experiment. Table

Table 1. The characteristics of the used Machines

Machine's name	Operating System	Processor	Memory
Dell-XPS L421X	Ubuntu (V.14.04 LTS)	1.8GHz*4 Intel Core i5	8 GB
Dell-Inspiron-3721	Ubuntu (V.14.04 LTS)	2.00GHz*4 Intel Core i5	4 GB
Dell-Inspiron-3521	Ubuntu (V.16.04 LTS)	1.8 GHz*4 Intel Core i5	6 GB
iMac-Early 2010	cinux Mint (V.17.1 Rebecca)	3.06GHz*2	4 GB
Dell-Inspiron-5559	Ubuntu (V.16.04 LTS)	2.30GHz*4 Intel Core i5	8 GB
iMac-Early 2009	OS X El Capitan (V.10.11.6)	2.93 *2 GHz Intel Core Due	8 GB
MacBook Air	OS X El Capitan (V.10.11.3)	1.6 *2 GHz Intel Core i5	8 GB

1 lists types of machines used to perform the experiments. The main goal here is to demonstrate the performance of the DDC in a heterogeneous distributed computing environment.

We used two benchmarks of datasets from Chameleon [20]. These are commonly used to test and evaluate clustering. Table 2 gives details about the datasets.

Table 2. Datasets

Benchmark	Size	Descriptions
D1	10,000 Points	Different shapes, with some clusters surrounded by others
D2	30,000 Points	2 small circles, 1 big circle and 2 linked ovals

The DDC approach is tested using various partitions of different sizes. Various scenarios were created based on the goals of the experiments. These scenarios mainly differ on the way the datasets are divided among the processing nodes of the distributed platform. For each scenario, we recorded the execution time for the local clustering, the merging step including contour calculations, aggregation time and Idle time.). finally we capture also the total execution time that the approach takes to finish all the steps. in the following we describe the different scenarios considered.

5.1 Experiment I

In this scenario we give each machine a random chunk of the dataset, the size of the partition that was generated for each machine is in the range between

1500 points and 10000 points. As the dataset is relatively small we chose eight machines for the computing platform.

Table 3 shows the execution time taken by each machine to run the algorithm (step one and step two) using synchronous and asynchronous communications respectively, it also shows the overall time taken to finish all the steps.

From Table 3, we can see that the time taken by each machine to accomplish the first step of the algorithm is the same for both synchronous and asynchronous, whereas the time of the second step is different. We can also notice that each machine returns different execution time of the whole algorithm. This is because the machines have different capacities (see Table 1).

The total execution time of the algorithm while using synchronous communication is smaller compared to when using asynchronous communication. This is because in synchronous communications, machines have more waiting time (up to 60% waiting time).

Table 3. Time (ms) taken by eight machines to run scenario I using synchronous and asynchronous communications

Machine	DS Size	Synchronous			Asynchronous		
		STEP01	STEP2	Time	STEP1	STEP2	Time
M1	10000	21270	1104	22374	21270	554	21824
M2	2500	1060	20862	21922	1060	2515	3575
M3	3275	5093	16930	22023	5093	2017	7110
M4	5000	4592	17644	22236	4591	2620	7211
M5	1666	227	21642	21869	227	391	618
M6	2000	292	21736	22028	292	416	708
M7	5000	7520	14665	22185	7515	13949	21464
M8	1500	200	21842	22042	195	4605	4800
Total Exec-Time				22374	Total Exec-Time		21824

5.2 Experiment II

In this scenario we allocate the whole dataset size to one machine and the remaining machines were allocated one eight of the dataset each. This scenario is chosen to show the worst case of waiting time.

Table 4 shows the execution time taken by each machine to execute the DDC technique (step one and step two) using synchronous and asynchronous communications respectively, it also shows the overall time taken to finish all the steps.

From Table 4, we can notice that the difference between the execution times of the synchronous and asynchronous DDC is still significant. Because with synchronous communications the machines need to wait for the last machine to finish its first step before they all start merging their results (step 2), whereas

for asynchronous model the seven machines did the merging (step2) while the last machine finishes its clustering (step1).

Table 4. Time (ms) taken by eight machines to run scenario II using synchronous and asynchronous communications

Machine	DS Size	Synchronous			Asynchronous		
		STEP1	STEP2	Time	STEP1	STEP2	Time
M1	10000	21270	973	22243	21270	595	21865
M2	1250	215	21775	21990	215	518	733
M3	1250	640	21383	22023	640	20100	20740
M4	1250	304	21730	22034	304	497	801
M5	1250	161	22034	22195	161	394	555
M6	1250	171	21856	22027	170	286	456
M7	1250	245	21918	22163	245	509	754
M8	1250	185	21854	22039	185	858	1043
Total Exec-Time				22243	Total Exec-Time	21865	

5.3 Experiment III

In this scenario we allocate to seven machines the whole dataset and the one machine was allocated one eight of the dataset. This scenario is chosen to show the effect of the complexity of the local clustering complexity on the machines and on the waiting time of some powerful machines.

Table 5 shows the execution time taken by each machine to run the algorithm (step one and step two) using synchronous and asynchronous communications respectively, it also shows the overall time taken to finish all the steps.

This scenario is the opposite of the previous scenario. Unlike the previous scenarios, Table 5 shows that the difference between the execution times of synchronous and asynchronous versions of the DDC is smaller. This is because in both cases the machines spend more time finishing the first step, therefore, the waiting time is less for synchronous over asynchronous model.

5.4 Experiment IV

In this scenario we took into account the machines capabilities and we divide the datasets according to their capacities. Therefore the work load is evenly distributed among them and we expect them to finish the first phase more or less at the same time. This allows to reduce the waiting time of the machines and follow immediately with the second phase. The total execution times of synchronous and asynchronous versions should be the same. This case favours more the synchronous implementation of the approach.

As predicted, Table 6 shows that there is no significant difference between the two execution times. Note that the little difference in favour of the synchronous

Table 5. Time (ms) taken by eight machines to run scenario III using synchronous and asynchronous communications

		Synchronous			Asynchronous		
Machine	DS Size	STEP1	STEP2	Time	STEP1	STEP2	Time
M1	10000	21270	35978	57248	21270	905	22175
M2	10000	21590	34869	56459	21590	11513	33103
M3	10000	53005	3008	56013	53005	3292	56297
M4	10000	32424	24691	57115	32424	6996	39420
M5	10000	17364	38493	55857	17364	4612	21976
M6	10000	15841	41237	57078	15841	2066	17907
M7	10000	38732	18483	57215	38727	18459	57186
M8	1250	185	56915	57100	184	16077	16261
		Total Exec-Time 57248			Total Exec-Time 57186		

version is due to the fact that in the asynchronous model the machines still need to execute the algorithm that checks which one finished first and receive the contours for merging.

Table 6. Time (ms) taken by eight machines to run scenario IV using synchronous and asynchronous communications

		Synchronous			Asynchronous		
Machine	DS Size	STEP1	STEP2	Time	STEP1	STEP2	Time
M1	1500	256	1505	1761	256	1159	1415
M2	1660	260	598	858	260	1512	1772
M3	500	252	1061	1313	252	626	878
M4	1000	253	621	874	253	608	861
M5	1500	255	1492	1747	255	600	855
M6	1400	260	605	865	260	514	774
M7	1000	259	1030	1289	259	939	1198
M8	1500	250	603	853	250	1500	1750
		Total Exec Time		1761	Total Exec Time		1772

5.5 Effective Speedup

The goal here is to compare our parallel clustering to the sequential algorithm and show the DDC speedup over the sequential version of clustering, as mentioned in Equation 3.

Considering the best scenario of executing the sequential version of DBSCAN on the fastest machine in the system. For instance, $T_1 = 15841$ ms. Clustering a partition of the same dataset on the same machine will take $= T_1^d = 258$ ms. The execution time of the DDC on the same datasets on eight heterogeneous machines with load balancing is $T_p = 1761$ ms (see Table 6). Therefore, from

equation 3, we can deduce a speedup of 9, which is still a super-linear speedup. In the next section we will show how many processing nodes are required to cluster a dataset of size N .

5.6 Scalability

The goal here is to show that the DDC technique scales well and also we can dynamically determine the optimal number of processing nodes required to cluster a dataset of size N . We consider two datasets, the first dataset D_1 contains 10,000 data points and the second D_1 contains 30,000 data points. Figure 4 shows the execution time (y_axis is in $\log_2$) against the number of machines in the system using the first dataset and Figure 5 shows the execution time (y_axis is in $\log_2$) against the number of machines in the system using the second dataset contains 30,000 data points.

Fig. 4. Scalability Experiment using dataset T_1 .

As one can see, from both Figures 4 and 5, the execution time of the first phase (Clustering and Contour) keeps decreasing as the number of machines in the distributed system increases. However, the time of the second phase (merging) keeps increasing gradually with the number of machines in the distributed system that is because the amount of communications in the second phase increases when the number of machines increases.

Fig. 5. Scalability Experiment using dataset T_2 .

In addition, the total execution time of the algorithm (which is the sum of the two times, phase one and two) keep decreasing as the number of processing nodes increases until it reaches a certain points where the total execution time starts to increase (at 8 machines for dataset D_1 and at 16 machines for dataset D_2). The optimal number of processing nodes required to execute DDC is returned when the overhead of the approach exceeds the execution time of the local clustering. This is a very interesting characteristic, as one can determine the number of machines that can be allocated in advance.

6 Conclusion

In this paper, we proposed an efficient and flexible distributed clustering framework that can work with existing data mining algorithms. The approach exploits the processing power of the distributed platform by maximising the parallelism and minimising the communications and mainly the size of the data that is exchanged between the nodes of the system. It is implemented using both synchronous and asynchronous communications, and the results were significantly in favour of the asynchronous model. The approach has an efficient data reduction phase which reduces significantly the size of the data exchanged therefore, it deals with the problem of communication overhead. The DDC approach has a super-linear speedup when the complexity of the local clustering has an NP

complexity. We also can determine the optimal number of processing nodes in advance.

Acknowledgement

The research work is conducted in the Insight Centre for Data Analytics, which is supported by Science Foundation Ireland under Grant Number SFI/12/RC/2289.

References

1. Aouad, L., Le-Khac, N.A., Kechadi, T.: Image analysis platform for data management in the meteorological domain. In: 7th Industrial Conference, ICDM 2007, Leipzig, Germany, July 14-18, 2007. Proceedings. vol. 4597, pp. 120–134. Springer Berlin Heidelberg (2007)
2. Aouad, L., Le-Khac, N.A., Kechadi, T.: Lightweight clustering technique for distributed data mining applications. *Advances in Data Mining. Theoretical Aspects and Applications* pp. 120–134 (2007)
3. Arlia, D., Coppola, M.: Experiments in parallel clustering with dbscan. In: European Conference on Parallel Processing. pp. 326–331. Springer (2001)
4. Bauer, E., Kohavi, R.: An empirical comparison of voting classification algorithms: Bagging, boosting, and variants. *Machine learning* 36(1), 105–139 (1999)
5. Bellifemine, F., Bergenti, F., Caire, G., Poggi, A.: Jadea java agent development framework. In: Multi-Agent Programming, pp. 125–147. Springer (2005)
6. Bendeche, M., Kechadi, M.T.: Distributed clustering algorithm for spatial data mining. In: Spatial Data Mining and Geographical Knowledge Services (ICSDM), 2015 2nd IEEE International Conference on. pp. 60–65. IEEE (2015)
7. Bendeche, M., Kechadi, M.T., Le-Khac, N.A.: Efficient large scale clustering based on data partitioning. In: Data Science and Advanced Analytics (DSAA), 2016 IEEE International Conference on. pp. 612–621. IEEE (2016)
8. Bendeche, M., Le-Khac, N.A., Kechadi, M.T.: Hierarchical aggregation approach for distributed clustering of spatial datasets. In: Data Mining Workshops (ICDMW), 2016 IEEE 16th International Conference on. pp. 1098–1103. IEEE (2016)
9. Brecheisen, S., Kriegel, H.P., Pfeifle, M.: Parallel density-based clustering of complex objects. *Advances in Knowledge Discovery and Data Mining* pp. 179–188 (2006)
10. Chaudhuri, A., Chaudhuri, B., Parui, S.: A novel approach to computation of the shape of a dot pattern and extraction of its perceptual border. *Computer vision and Image Understanding* 68, 257–275 (03 December 1997)
11. Chen, M., Gao, X., Li, H.: Parallel dbscan with priority r-tree. In: Information Management and Engineering (ICIME), 2010 The 2nd IEEE International Conference on. pp. 508–511. IEEE (2010)
12. Coppola, M., Vanneschi, M.: High-performance data mining with skeleton-based structured parallel programming. *Parallel Computing* 28(5), 793–813 (2002)
13. Cortese, E.: Benchmark on jade message transport system. URL: <http://jade.cselt.it/doc/tutorials/benchmark/JADERTTBenchmark.htm> (2005)
14. Dhillon, I., Modha, D.: A data-clustering algorithm on distributed memory multiprocessor. In: large-Scale Parallel Data Mining, Workshop on Large-Scale Parallel KDD Systems, SIGKDD. pp. 245–260. Springer-Verlag London, UK (1999)

15. Dhillon, I.D., Modha, D.S.: A data-clustering algorithm on distributed memory multiprocessors. In: *Large-Scale Parallel Data Mining*. pp. 245–260. Springer Berlin Heidelberg (2000)
16. Duckhama, M., Kulikb, L., Worboysc, M., Galtond, A.: Efficient generation of simple polygons for characterizing the shape of a set of points in the plane. Elsevier Science Inc. New York, NY, USA 41, 3224–3236 (15 March 2008)
17. Edelsbrunner, H., Kirkpatrick, D.G., Seidel, R.: On the shape of a set of points in the plane. *Information Theory, IEEE Transactions on* 29(4), 551–559 (1983)
18. Ester, M., Kriegel, H.P., Sander, J., Xu, X.: A density-based algorithm for discovering clusters in large spatial databases with noise. In: *Kdd*. vol. 96, pp. 226–231 (1996)
19. Fadilia, M., Melkemib, M., ElMoataza, A.: *Pattern Recognition Letters: Non-convex onion-peeling using a shape hull algorithm*, vol. 24. ELSEVIER (15 October 2004)
20. Fränti, P.: *Clustering datasets* (2015), <http://cs.uef.fi/sipu/datasets/>
21. Fu, Y.X., Zhao, W.Z., Ma, H.F.: Research on parallel dbscan algorithm design based on mapreduce. In: *Advanced Materials Research*. vol. 301, pp. 1133–1138. Trans Tech Publ (2011)
22. Garg, A., Mangla, A., Bhatnagar, V., Gupta, N.: Pbirch: A scalable parallel clustering algorithm for incremental data. *10th Int'l. Symposium on Database Engineering and Applications (IDEAS-06)* pp. 315–316 (2006)
23. Geng, H., Deng, X., Ali, H.: A new clustering algorithm using message passing and its applications in analyzing microarray data. In: *Machine Learning and Applications, 2005. Proceedings. Fourth International Conference on*. pp. 6–pp. IEEE (2005)
24. Ghosh, S.: *Distributed systems: an algorithmic approach*. CRC press (2014)
25. Guo, Y., Grossman, R.: A fast parallel clustering algorithm for large spatial databases, high performance data mining. *Data Mining and Knowledge Discovery* (2002)
26. Han, J., Pei, J., Kamber, M.: *Data mining: concepts and techniques*. Elsevier (2011)
27. Jain, A.K., Murty, M.N., Flynn, P.J.: Data clustering: a review. *ACM computing surveys (CSUR)* 31(3), 264–323 (1999)
28. Januzaj, E., Kriegel, H.P., Pfeifle, M.: *Advances in Database Technology - EDBT 2004: 9th International Conference on Extending Database Technology*, chap. DBDC: Density Based Distributed Clustering, pp. 88–105. Springer Berlin Heidelberg (2004)
29. Laloux, J.F., Le-Khac, N.A., Kechadi, M.T.: Efficient distributed approach for density-based clustering. *Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE), 20th IEEE International Workshops* pp. 145–150 (27-29 June 2011)
30. Le-Khac, N.A., Bue, M., Whelan, M., M-T.Kechadi: A knowledgebased data reduction for very large spatio-temporal datasets. *International Conference on Advanced Data Mining and Applications, (ADMA2010)* (19-21 November 2010)
31. Melkemi, M., Djebali, M.: Computing the shape of a planar points set. *Elsevier Science* 33, 14231436 (9 September 2000)
32. Moreira, A., Santos, M.Y.: Concave hull: A k-nearest neighbours approach for the computation of the region occupied by a set of points. In: *Int'l. Conf. on Computer Graphics Theory and Applications (GRAPP-2007)*. pp. 61–68. Barcelona, Spain (8-11 March 2007)
33. Rokach, L., Schclar, A., Itach, E.: Ensemble methods for multi-label classification. *Expert Systems with Applications* 41, 7507 – 7523 (2014)

34. Solar, R., Borges, F., Suppi, R., Luque, E.: Improving communication patterns for distributed cluster-based individual-oriented fish school simulations. *Procedia Computer Science* 18, 702–711 (2013)
35. Wu, X., Zhu, X., Wu, G.Q., Ding, W.: Data mining with big data. *IEEE Transactions on Knowledge and Data Engineering* 26(1), 97–107 (2014)
36. Xu, X., Jger, J., Kriegel, H.P.: A fast parallel clustering algorithm for large spatial databases. *Data Mining and Knowledge Discovery archive* 3, 263–290 (September 1999)
37. Zaki, M.J., Ho, C.T.: Large-Scale Parallel Data Mining, chap. Parallel and Distributed Data Mining: An Introduction, pp. 1–23. Springer Berlin Heidelberg, Berlin, Heidelberg (2000)
38. Zhang, T., Ramakrishnan, R., Livny, M.: Birch: an efficient data clustering method for very large databases. In: *ACM Sigmod Record*. vol. 25, pp. 103–114. ACM (1996)
39. Zhou, A., Zhou, S., Cao, J., Fan, Y., Hu, Y.: Approaches for scaling dbscan algorithm to large spatial databases. *Journal of computer science and technology* 15(6), 509–526 (2000)