

μ μ ,
μ , μ
μ . μ
μ . μ
μ , μ
μ μ μ JAVA
μ μ
μ μ
μ , μ
μ μ μ ,
μ

Contents

.....	1
μ	2
.....	3
.....	4
Abstract	5
1.	7
2. Hadoop	9
2.1 μμ MapReduce.....	10
2.2 μ Hadoop	11
2.3 Hadoop Common	12
2.4 Hadoop Yarn	12
3.	13
4. HDD SSD.....	17
4.1 read write	17
4.2 read write Hadoop ...	18
5. Density-Based Algorithms	22
6. μ μ μ μ	25
6.1 DB-CURE MapReduce	25
6.2 DBSCAN-MR	27
6.3 CFSFDP-MR (Clustering by fast search and find of density Peaks).....	30
6.4 Transitive heuristic-MR.....	31
7. μ μ	32
7.1 Dataset μ	32
7.2 μ	33
8. μ μ μ	38

1:	μ	Hadoop.(Source: [2]).....	10
2:		μ MapReduce (Source:[3])	11
3:		(Source : [13])	18
4:		μ (Source:[14])	19
5:	HDD	SSD MapReduce tasks μ normalize μ	
		(Source:[15])	20
6:		. (Source : [15]).....	21
7:		μ (Shuffle).	
		(Source: [15])	21
8:	μ	DBCURE HDD SSD.....	34
9:	μ	DBSCAN HDD SSD.....	34
10:	μ	CFSFDP HDD SSD.....	35
11:		μ μ HDD.....	35
12:		μ μ SSD.....	35
13:	μ	Transitive Heuristic HDD SSD.....	36

Περίληψη

MapReduce μ μ μ μ μ
μ μ μ μ μ
(divide and conquer) μ , μ
, μ μ . Hadoop (Yahoo) μ open source μ
MapReduce Java. μ , μ
Hadoop μ
Hard Disk Drive Solid State Disk.
μ , μ μ μ
μ μ , μ framework Apache Hadoop. H
μ . μ density-based
clustering μ μ
μ μ Hard Disk Drive (HDD) Solid
State Disks (SSD). μ μμ
JAVA μ μ μ
μ μ (Virtual Machine – VM) .
μ μ μ μ μ
, μ μ μ μ μ
μμ μ MapReduce computer cluster.

Abstract

MapReduce is one of the most popular distributed computing platforms for applications that process largescale data. It has been applied to many areas of divide and conquer problems such as search engines, data mining and indexing of data. Hadoop, initially developed by Yahoo, is an open source software with the ability of implementing the MapReduce model in Java environment. This Dissertation is focused on the examination of Hadoop's performance when operating in Hard Disks (HDDs) and Solid-State Disks (SSDs) disks. More specifically, the dissertation aims at studying the performance of Density-based algorithms implemented within Apache Hadoop. The present study can be divided in two parts. In the first part, the implementation of density based algorithmic methods is explored while in the second part, the differences that arise HDD and SSD are quantified through numerical tests. Implementation of the algorithms are utilized in JAVA programming language and the tests are performed using a Virtual Machine (VM) built for this purpose. The ultimate goal of this diploma thesis is to draw conclusions about the duration of executing parallel, distributed algorithm groups using the MapReduce programming model in a cluster of computers.

1. Εισαγωγή

Η τεχνολογία των μεγάλων δεδομένων (Big Data) αποτελεί ένα σημαντικό πεδίο έρευνας και ανάπτυξης στην πληροφορική (Information Technology - IT). Τα τελευταία χρόνια, η ποσότητα των δεδομένων που παράγονται και συλλέγονται έχει αυξηθεί εκθετικά, κυρίως λόγω της ανάπτυξης των κοινωνικών μέσων επικοινωνίας, των υπηρεσιών διαδικτύου και των εφαρμογών που απαιτούν την επεξεργασία μεγάλων όγκων πληροφοριών. Σύμφωνα με την International Data Corporation (IDC), η ποσότητα των δεδομένων που παράγονται παγκοσμίως θα φτάσει τα 160 zettabytes (ZB) έως το 2025.

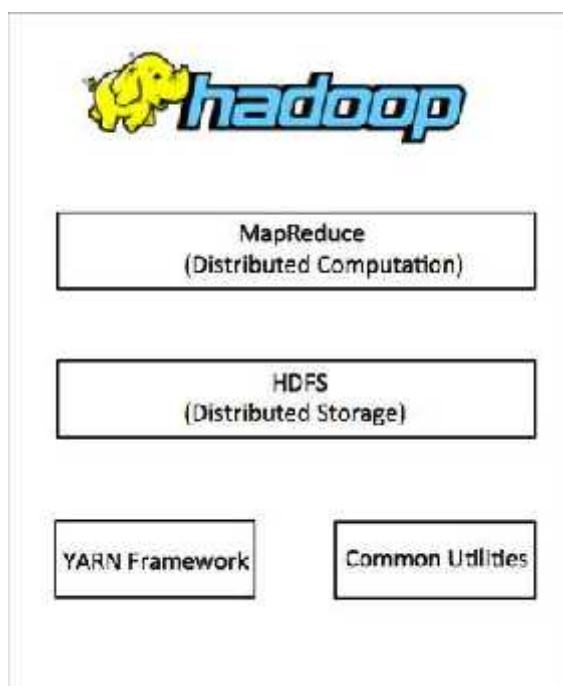
Η τεχνολογία των μεγάλων δεδομένων (Big Data) απαιτεί την ανάπτυξη νέων εργαλείων και μεθόδων για την αποθήκευση, την επεξεργασία και την ανάλυση των δεδομένων. Τα εργαλεία αυτά πρέπει να είναι σε θέση να χειρίζονται την τεράστια ποσότητα των δεδομένων, να παρέχουν υψηλή απόδοση και να είναι εύκολα στην χρήση. Τα πιο δημοφιλή εργαλεία για την επεξεργασία μεγάλων δεδομένων είναι τα frameworks όπως το MapReduce, το Spark, το Hadoop και το Storm.

Τα εργαλεία αυτά βασίζονται στην αρχιτεκτονική των μεγάλων δεδομένων, η οποία είναι βασισμένη στην κατανομή των δεδομένων σε έναν μεγάλο αριθμό υπολογιστών (computer cluster). Η κατανομή αυτή επιτρέπει την επεξεργασία των δεδομένων παράλληλα, γεγονός που οδηγεί σε σημαντική αύξηση της απόδοσης. Η αρχιτεκτονική των μεγάλων δεδομένων απαιτεί επίσης την ύπαρξη ενός αποτελεσματικού συστήματος αποθήκευσης, το οποίο μπορεί να χειριστεί την τεράστια ποσότητα των δεδομένων. Τα συστήματα αποθήκευσης που χρησιμοποιούνται για την αποθήκευση μεγάλων δεδομένων είναι τα Hard Disk και τα Solid State.

2. Hadoop


Hadoop is a distributed system consisting of a master node and several slave nodes. The master node is responsible for managing the cluster and the slave nodes are responsible for storing and processing the data. Hadoop is built on top of the Linux operating system and uses Java as the programming language. The core of Hadoop is the MapReduce framework, which allows for the parallel processing of large datasets. Hadoop also includes the Hadoop Distributed File System (HDFS), which is a distributed file system that allows for the storage of large amounts of data across multiple nodes. Hadoop is used in a wide variety of applications, including data analysis, machine learning, and scientific computing. Hadoop is an open source project and is licensed under the Apache License. Hadoop is a key component of the Big Data ecosystem.

1. Hadoop (Hadoop Distributed File System-HDFS).
2. “MapReduce” framework. The “map” function is used to process the input data and the “reduce” function is used to aggregate the results. Hadoop uses a distributed file system (HDFS) to store the data. Hadoop is written in Java.
3. “Hadoop Common”, which is the core of Hadoop. Hadoop is written in Java.
4. Hadoop. Hadoop is written in Java. Hadoop is a key component of the Big Data ecosystem.




1: μ Hadoop.(Source: [2])

2.1 Προγραμματιστικό Μοντέλο MapReduce



reduce. reduce μ map, μ μ
 <key3,value3> μ μ μ key.
 reduce .



2: μ MapReduce (Source:[3])

MapReduce μ master-slave. μ
 (NameNode)
 slaves. μ , NameNode (job
 tracker) slaves (DataNodes)
 (task trackers) . , task trackers job
 tracker [4].

2.2 Το σύστημα αρχείων του Hadoop

μ μ Hadoop (HDFS)
 μ μ
 μ μ . HDFS μ
 μ μ μ μ μ μ
 μ . μ MapReduce, μ
 μ μ cluster, μ
 μ μ μ . μ HDFS
 μ . μ μ μ μ
 Hadoop μ
 μ μ [1, 5]. Hadoop
 Distributed File System (HDFS) μ “Write once-Ready many”.
 , μ
 .
 μ master (NameNode)
 μ μ data blocks
 (DataNodes) μ

2.3 Hadoop Common

Hadoop Common, (Java)
(Windows, Unix
Hadoop.

2.4 Hadoop Yarn

Hadoop Hadoop Yarn.
Hadoop cluster. Yarn
Resource Manager
Application Master,
(Job).
Manager physical plan.
[6].

4. HDD SSD

μ (HDD).
 , 40% ,
 / (Input / Output) μ 2% . μ
 μ μ , Web, cloud μ
 virtualization ,
 μ μ μ , μ μ μ
 μ GB. 2011, μ
 , SSD, . SSD ,
 μ .
 μ .
 μ μ μ μ μ
 μ μ μ , μ μ , μ
 μ μ μ μ μ
 μ SSD μ μ μ
 μ μ access latency input/output
 (IOPS) μ [11].

4.1 read write

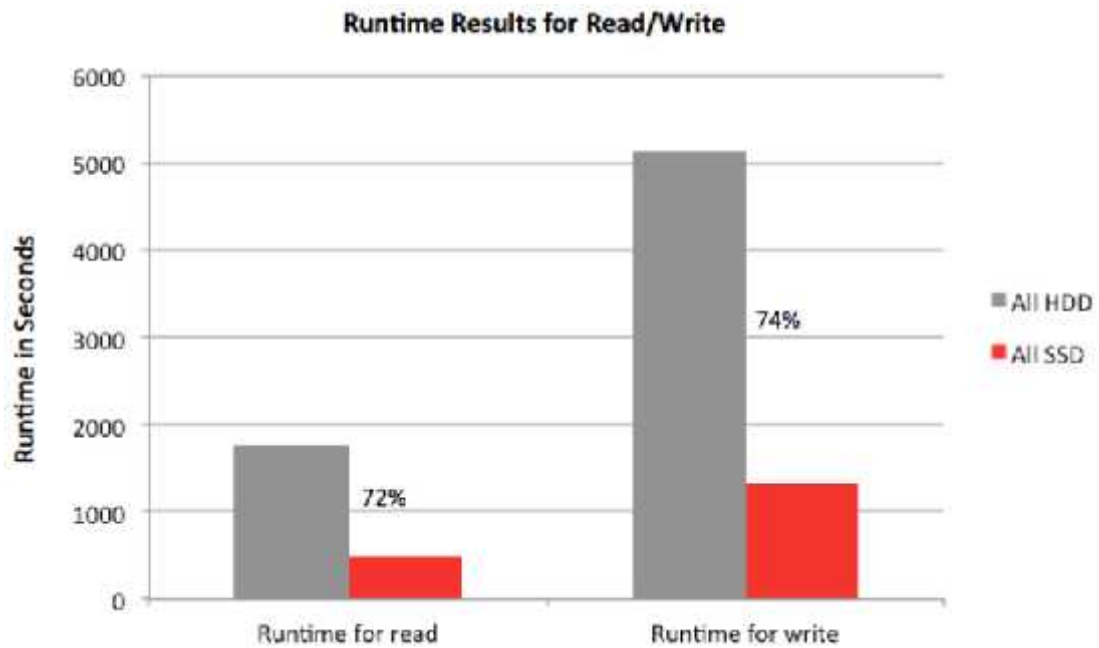
SSD / μ 500Mb / sec.
 μ 100Mb / sec HDD,
 μ μ μ 5
 , μ . μ
 μ μ
 SSD , μ
 SSD μ HDD.
 / μ μ , SSD
 μ μ HDD μ
 μ μ . μ μ μ
 , μ μ SSD
 μ HDD [12].

4.2

read write

Hadoop

HDD SSD
 Hadoop Distributed File System.
 70%-75%
 Hadoop [13].

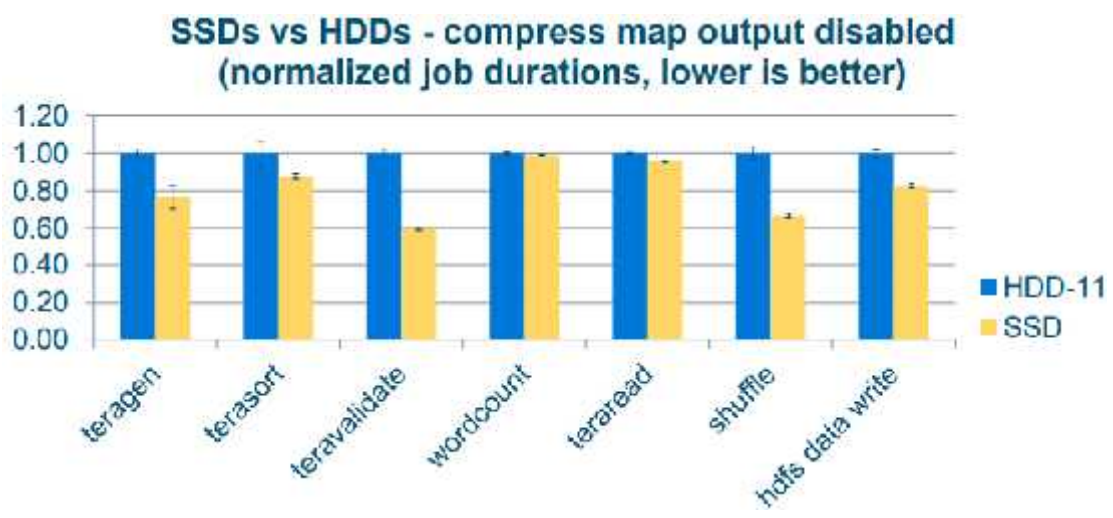


3: (Source : [13])

HDD SSD.
 megabyte (Mb/s throughput)
 (file size)
 SSD (SSD2) 120 GB (read) SSD
 HDD 500 GB HDD 1 TB.
 MB/s
 HDD (SSD1)
 SSD (SSD2)

μ , μ
 SSD μ μ HDD
 μ (Teragen) , μ (Terasort) ,
 μ μ μ (WordCount),
 μ (Shuffle) μ
 map reduce Hadoop
 μ μ (Teravalidate). SSD
 μ HDD

[15].

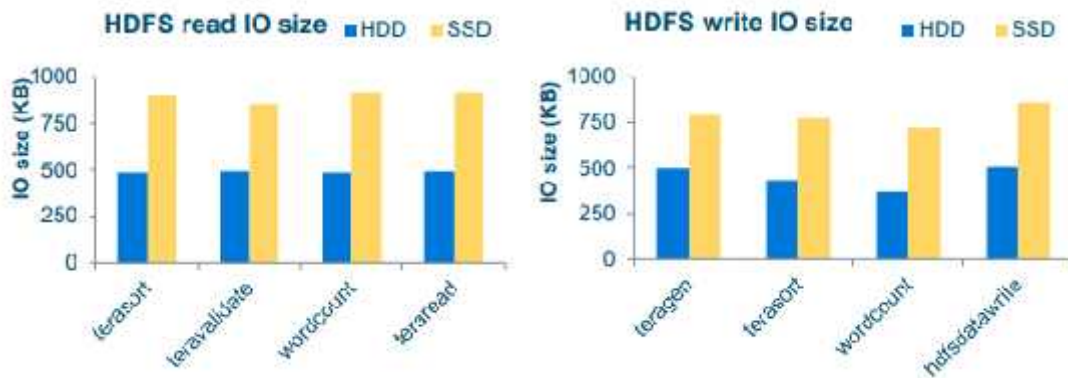


5: HDD SSD MapReduce tasks μ normalize μ
 (Source:[15])
 μ normalize. μ
 SSD μ 10% (Teraread)
 40% (Teravalidate). μ
 WordCount.

μ
 μ
 SSD μ μ
 KB μ

HDD.

Reason 1: SSDs > HDDs for seq IO size



6:

(Source : [15])

μ

Shuffle

μ

μ

map

reduce.

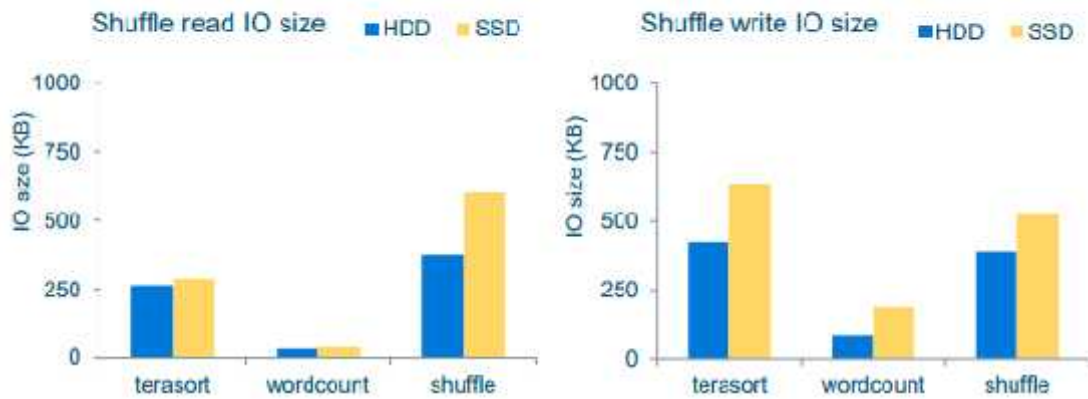
WordCount,

SSD

μ

HDD.

Reason 2: SSDs > HDDs for small IO in shuffle



7:

(Source: [15])

μ

(Shuffle).

5. Density-Based Algorithms

- DBSCAN [16].
1. Compute the distance matrix D .
 2. For each point p in the dataset, find its ϵ -neighborhood $N_\epsilon(p)$ and count the number of points in $N_\epsilon(p)$ that are not p (core).
 3. If the number of points in $N_\epsilon(p)$ is greater than or equal to minPts , then p is a core point.
 4. If p is a core point, then p and all points in $N_\epsilon(p)$ that are not core points form a cluster.
 5. Repeat steps 2-4 for all points in the dataset.
- (noise).

- Clustering by fast search and find of density peaks (dc, minPts) [17].
1. Compute the distance matrix D .
 2. For each point p in the dataset, find its ϵ -neighborhood $N_\epsilon(p)$ and count the number of points in $N_\epsilon(p)$ that are not p (cut-off distance – dc).
 3. If the number of points in $N_\epsilon(p)$ is greater than or equal to minPts , then p is a density peak.

Dataset

4. μ (cluster centers)

5. Cluster centers (cluster).

DBCURE (,) [18].

1. μ D
2. μ (μ)
3. μ S.
4. μ 3 μ μ S μ μ .

Transitive heuristic [19]

1. Dataset
2. Condorcet (attribute μ μ cluster μ)
3. cluster μ μ cluster
4. cluster center).

5. μ μ . μ μ , μ 1.

6.

$\mu \mu$

μ

μ

$\mu \mu$ MapReduce $\mu \mu \mu$.

6.1 DB-CURE MapReduce

μ DB-CURE MR

1. μ μ (covariance matrix)

2. $\mu \mu$: μ , μ .

3. $\mu \mu$. $\mu \mu$ (core points): core point p, density reachable μ . $\mu \mu$ μ . $\mu \mu$ μ μ μ μ . $\mu \mu$ μ . $\mu \mu$ μ μ , $\mu \mu$ μ μ .

4. core points $\mu \mu \mu$ (core points): $\mu \mu \mu \mu$. $\mu \mu \mu$.

μ paper μ $\mu \mu \mu$ μ -Tree. μ clusters μ global.

MapReduce. μ (Dataset) D μ (cells) $\mu \mu \mu$. $\mu \mu \mu \mu$ $\mu \mu$ (cell-ID)- $\mu \mu$. μ Si μ i- μ G_{21_kd} S_{61} μ μ μ $i j$ μ G_{21_kd} . :

1. π μ ID g map $\langle g, \pi \rangle$. reduce

$\langle g, ListOfItems\langle g \rangle \rangle$.

2.

map <g,ListOfItems<g>>
 reduce 1 μ S_{xx}S_{yy}S_xS_y μ <g', n,
 <S_{xx}S_{yy}S_xS_y> <g,ListOfItems> , g' g. reduce
 μ μ μ ID (g) μ
 <g, Σn, list <S_{xx}S_{yy}S_xS_y>, ListOfItems<g>>.

3.

μ map
 <g, n list <S_{xx}S_{yy}S_xS_y>, ListOfItems<g>>
 reduce 2 μ
 g.
 μ μ :

$$\Sigma(u, v) = \frac{1}{n-1} \sum S_{uv}[k_1, \dots, k_d] - \frac{n}{n-1} \mu(u)\mu(v)$$

n μ
 2 μ_i(u) = $\frac{1}{n} \sum [S_{ui}[k_1], \dots, k_d]$ μ μ
 μ u μ . map
 (key) <cell-ID, μ > μ
 (value) μ μ . reduce μ
 μ μ μ (clusters).

4.

map μ reduce 3
 μ μ . μ
 μ μ μ μ μ
 μ μ μ μ μ
 μ μ μ (cluster). μ μ
 μ μ μ μ . μ
 μ cluster. μ cluster

cluster . μ
 cluster μ μ cluster
 μ . cluster cell.
 μ cluster
 map μ
 cluster μ μ μ
 μ μ .

$$\left[p_i(l) - \sqrt{c * \Sigma(l,l)}, p_i(l) + \sqrt{c * \Sigma(l,l)} \right]$$

$c = -2 \log(\tau(2\pi)^{d/2} | \Sigma |^{1/2})$. μ μ Gaussian
 p μ μ
 (μ μ)
 $f(p_i) \geq \tau$. reduce μ
 μ <key, list(clusters)>.

5.

map reduce μ μ
 μ clusters , μ
 clusters μ . ,
 μ cluster μ μ
 μ clusters . μ
 μ map μ cluster
 (cluster-ID) μ μ μ μ cluster.
 reduce.

6.2 DBSCAN-MR

μ paper
 μ cluster
 R-TREE μ μ μ μ
 μ μ (3 μ), μ μ
 μ . μ
 μ (Dataset) D μ D μ
 μ μ μ μ μ μ
 μ μ μ μ μ μ

μ μ (cell) μ
(ID). μ μ
 cell ID. μ 3 .

1.

map μ μ μ
 (ID) . (key) μ (ID)
 value μ μ . **reduce** μ μ μ
 μ μ
 μ $\langle \text{ID}, \text{list}(\text{values}(\text{key})) \rangle$.

2.

map **reduce** 1
 μ μ (dc) μ
 (minPts) μ μ μ .
 μ μ μ
 μ μ dc. μ μ minPts
 μ μ (core point cluster center) μ
 μ dc . μ μ
 μ μ (cluster). μ μ cluster.
 density-reachable cluster. μ , cluster
 μ μ ,
 μ cluster.
 cluster μ μ μ
 μ . μ μ (clusters) key
 value μ μ μ
reduce μ μ μ clusters
 μ map μ clusters

3.

map μ $\langle \text{key}, \text{list}(\text{clusters}) \rangle$. μ
 dc core points density-reachable
 μ clusters .
 μ μ μ μ cluster (cluster ID) μ
 key map. map cluster μ
 μ $\langle \text{clusterID}, \text{ListofItems} \rangle$. **reduce.**

6.3 CFSFDP-MR (Clustering by fast search and find of density Peaks)

μ 3 .

1.

map μ μ μ .

μ μ μ

(dc) - μ 1- μ .

map (key) μ ()

μ (value) μ , μ

reduce identity reducer μ

μ .

2.

map μ μ

μ μ .

(key) μ (μ

μ μ) μ (value) μ < , ,

μ > . **reduce** μ μ

(key) μ μ map 2.

3.

map 2

μ μ μ 2

. μ

cluster center μ μ

(nose). cluster center μ μ μ

. cluster center

cluster μ μ

μ μ

cluster. **reduce** μ key (Cluster

Number) <key,list(values)>.

6.4 Transitive heuristic-MR

μ μ . **map**
 μ n attributes . Condorcet
 cluster center μ μ .
 μ cluster center μ attributes μ . map
 cluster center attributes , value attributes
 .
 Condorcet: **map** μ μ
 μ μ . μ cluster centers
 μ Condorcet. μ μ , attributes
 μ cluster centers. attributes μ
 , μ . , μ μ
 μ μ μ cluster
 centers . μ μ
 μ μ μ cluster.
 μ μ μ μ cluster attributes
 cluster center μ . **reduce** μ μ
 μ cluster attributes cluster attributes μ (value)
 μ cluster, list(elements).

7. μ μ μ

μ μ μ μ μ μ μ μ μ μ

μ HDD SSD. μ μ μ μ μ μ μ

μ μ μ μ μ μ μ μ μ μ

7.1 Dataset μ

1: .

4 CPU	Intel® Xeon® W3550 @3.07 GHz
RAM	4 GB
DISK 1(HDD)	100GB Samsung 850 Evo
DISK 2(SSD)	100GB Seagate ST31000524AS

2: μ μ

OS	Centos 7
Java SDK	Oracle Java 1.8.0
Hadoop version	Hadoop-3.1

3 μ Dataset

Dataset 1	Size : 500 items	Number of Attributes : 2
Dataset 2	Size : 1500 items	Number of Attributes: 2
Dataset 3	Size: 4000 items	Number of Attributes: 2
Dataset	μ make_blobs	python.
μ μ μ μ μ μ μ μ μ μ		.

μ	n	μ	clustering.
size = 500 items	μ	n=4	μ
size = 1500 items	μ	μ	n=6. μ
μ		size = 4000	μ n=10 μ
μ	μ	μ	μ
Dataset 4	Size: 22000 items	Number of Attributes: 2	
Dataset 5	Size: 30000 items	Number of Attributes: 2	
Dataset	μ	Birch Dataset	μ online.
(Source : [20]).		μ	Dataset
clustering μ	μ	.	

	4	μ	
Dataset 1	Size: 8124 items	Number of Attributes: 22	
Dataset 2	Size: 15000items	Number of Attributes: 22	
Dataset 3	Size: 22000items	Number of Attributes: 22	
Heuristic-MR	μ	μ	Transitive
μ [21].	μ	Dataset Mushroom	
μ μ	μ	Dataset	Number of Attributes = 22.
μ	μ	μ	python
μ	μ	μ	Attribute.

7.2

μ μ μ

μ μ μ .

DBCURE μ μ μ

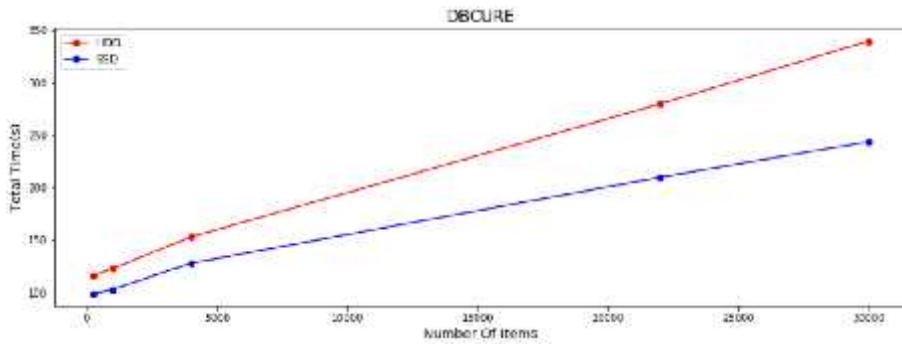
μμ μ μ μ .

5 μ μ μ ,

μ μ μ . μ μ ,

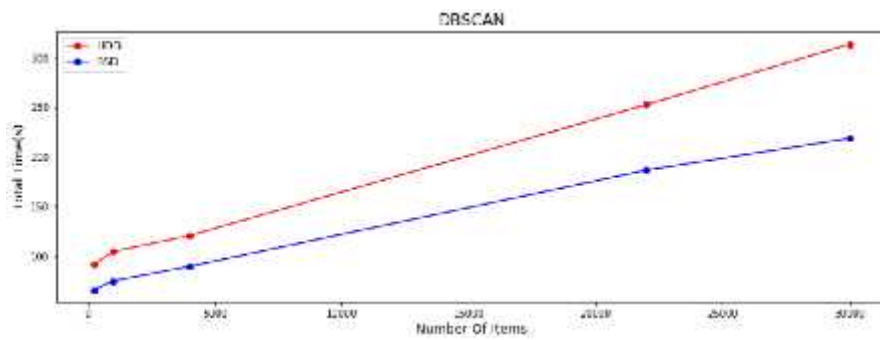
μ μ , DBCURE SSD 15% 30%

HDD.



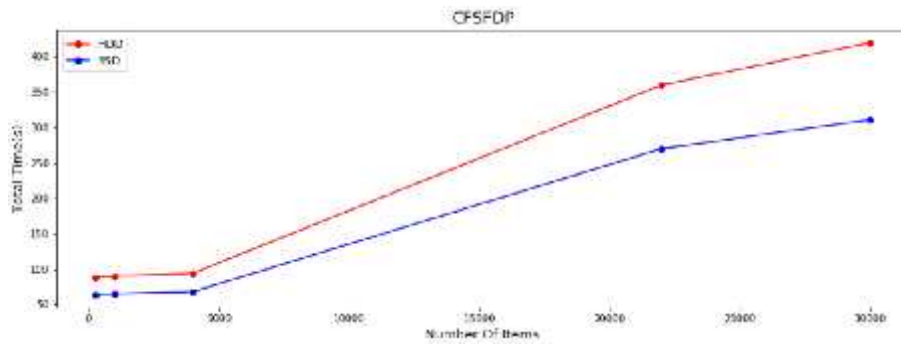
8: μ DBCURE HDD SSD.

μ DBSCAN. μ μ DBCURE,
 DBSCAN μ
 HDD SSD. μ , μ .
 μ 25% 30%. μ ,
 μ μ μ μ μ .



9: μ DBSCAN HDD SSD.

, μ μ μ Clustering
 by fast search and find of density Peaks. μ .
 μ μ μ μ ,
 μ μ μ .
 . , μ μ μ
 26% μ 28% μ HDD SSD.



10: μ CFSFDP HDD SSD.

μ

μ

μ

$\mu\mu$

μ

μ

μ

μ

CFSFDP,

DBCURE.

μ

DBSCAN

DBCURE.

μ

CFSFDP

μ

μ

DBCURE

μ DBSCAN (

μ

μ

.

DBSCAN

μ

μ

μ

μ

μ

μ

CFSFDP

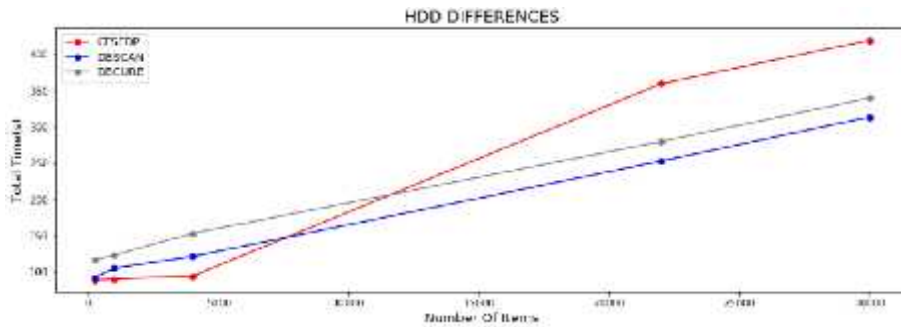
μ

μ

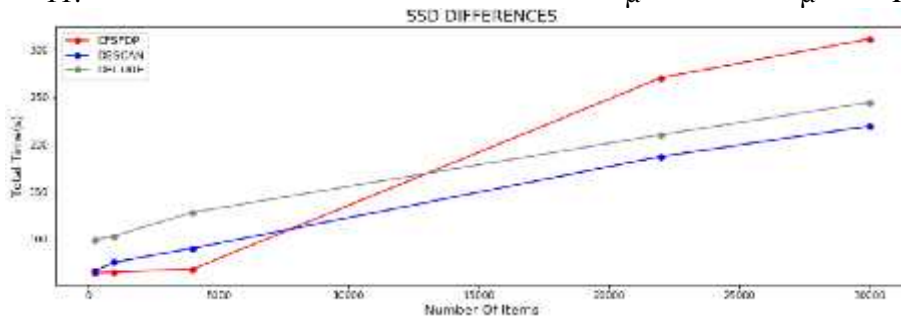
μ

μ

.

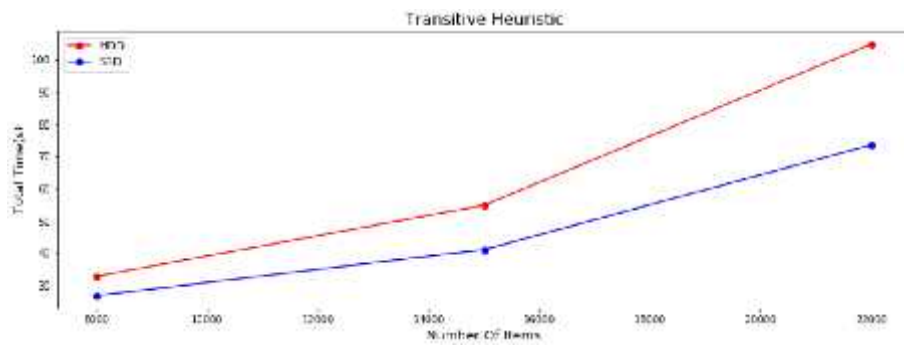


11: μ μ HDD.



12: μ μ SSD.

μ Transitive Heuristic. μ HDD
 μ SSD ,
 μ SSD HDD.
 μ μ . μ



13: μ Transitive Heuristic HDD SSD.

3. μ Hadoop cluster μ μ μ
 μ μ .

Βιβλιογραφία

1. Ghazi, M.R. and D. Gangodkar, *Hadoop, MapReduce and HDFS: A Developer's Perspective*, in *International Conference on Intelligent Computing, Communication & Convergence*. 2015, Elsevier: Bhubaneswar, Odisha, India. p. 45-50.
2. <https://www.tutorialspoint.com/>.
3. www.storegaga.com.
4. Dean, J. and S. Ghemawat, *MapReduce: Simplified Data Processing on Large Clusters*. ACM, 2008. **51**(1): p. 107-113
5. Shvachko, K., et al., *The Hadoop Distributed File System in 2010 IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST)*. 2010, IEEE: Incline Village, NV, USA. p. 1-10
6. Vavilapalli, V.K., et al., *Apache Hadoop YARN: Yet Another Resource Negotiator*, in *4th annual Symposium on Cloud Computing*. 2013, ACM: Santa Clara, California.
7. Ghuman, S.S., *Clustering Techniques- A Review* International Journal of Computer Science and Mobile Computing, 2016. **Vol. 5**(5): p. 524-530.
8. Garima, H. Gulati, and P.K.Singh, *Clustering Techniques in Data Mining: A Comparison in 2nd International Conference on Computing for Sustainable Global Development*. 2015, : New Delhi, India.
9. Jain, A.K. and R.C. Dubes, *Algorithms for clustering data*. 1988: Prentice Hall PTR.
10. Gandhi, G. and R. Srivastava, *Review Paper: A Comparative Study on Partitioning Techniques of Clustering Algorithms*. IJCA, 2014. **87**: p. 10-13.
11. Kasavajhala, V., *Solid State Drive vs. Hard Disk Drive Price and Performance Study*. 2011: Dell PowerVault Storage Systems.
12. www.java-performance.info.
13. <https://blog.westerndigital.com>.
14. M.Bakratsas, et al., *Hadoop MapReduce Performance on SSDs for Analyzing Social Networks*. Big Data Research, 2017. **11**: p. 1-10.
15. Kambatla, K. and Y. Chen, *The Truth About MapReduce Performance on SSDs*, in *28th Large Installation System Administration Conference*. 2014: Seattle, WA.
16. He, Y., et al., *MR-DBSCAN: a scalable MapReduce-based DBSCAN algorithm for heavily skewed data*. Frontiers of Computer Science, 2014. **8**(1): p. 83-99.
17. Pang, L. and F.A. Liu, *Density-Based Algorithm in MapReduce*, in *2016 9th International Symposium on Computational Intelligence and Design (ISCID)*. 2016. p. 394-397.
18. Kim, Y., et al., *DBCURE-MR: An efficient density-based clustering algorithm for large data using MapReduce*. Information Systems 2013. **42**: p. 15-35.
19. Lamari, Y. and S.C. Slaoui, *Clustering categorical data based on the relational analysis approach and MapReduce*. Journal of Big Data, 2017. **4**(1): p. 28.
20. <https://cs.joensuu.fi/sipu/datasets/>.
21. <https://archive.ics.uci.edu/ml/datasets.html>.