



وقائع المؤتمر العلمي البحثي الدوري الثامن للباحثين من حملة الشهادات العليا  
شعبة البحوث والدراسات التربوية / قسم الاعداد والتدريب وبالتعاون مع مركز  
البحوث والدراسات التربوية / وزارة التربية وجامعة بغداد / كلية التربية ابن رشد  
والجامعة المستنصرية – كلية التربية الاساسية والمنعقد تحت شعار  
((الاستدامة ودورها في تنمية القطاع التربوي))

للمدة 2025/2/12

## Optimizing Big Data Processing: An Analysis of Map Reduce Scheduling Algorithms and Frameworks

**Mustafa Razaq Kalaf<sup>(1)</sup>**  
Ministry of Education, Baghdad, Iraq  
07763162588

**Mohammed Rashid Dubayan<sup>(2)</sup>**  
Imam Al kadhum college(IKC), Baghdad, Iraq  
07715231156

[mostafarazaq98@gmail.com](mailto:mostafarazaq98@gmail.com)

[mohamed313rashid@gmail.com](mailto:mohamed313rashid@gmail.com)

### Abstract:

Google makes use of the MapReduce programming model in order to process enormous amounts of data in an effective manner inside a distributed computing environment. Typically, it is employed for the purpose of doing distributed computing on clusters of computers. The computational processing of data is often performed on data that is stored in either a file system or a database. The MapReduce framework leverages the principle of data localization, enabling the processing of data in close proximity to the storage locations. This approach effectively mitigates the need for superfluous data transport. Recent developments in the field of big data have revealed a notable surge in the volume of data, demonstrating an exponential growth pattern. In recent years, this phenomenon has served as a source of inspiration for numerous scholars, prompting them to delve into novel avenues of inquiry within the realm of big data across various domains. The widespread appeal of large data processing platforms that make use of the MapReduce architecture is a primary factor driving the growing interest in improving the platforms' performance. The improvement of resources and job scheduling is of utmost importance as it plays a crucial role in determining whether applications can successfully meet performance objectives in various usage scenarios. The significance of scheduling in big data is significant, primarily focused on the optimization of execution time and cost associated with processing. The purpose of this research is performing an examination of scheduling in

MapReduce, with a particular emphasis on two important aspects: Nomenclature as well as operational efficiency analysis.

**Keywords:** MapReduce, Big Data, Scheduling Algorithms, Hadoop, Resource Optimization, Distributed Systems

**Note:** The research is based on an M.A thesis or a PhD dissertation (if any).NO

### Introduction:

In the present period, a substantial volume of data is being generated from many sources such as scientific instruments, digital media, and web publishing, among others. The efficient storage, retrieval, and analysis of this data has emerged as a significant challenge for the computing industry. To address the substantial storage and processing requirements associated with vast quantities of data, entails the construction of a computer system at the scale of a data centre.

1- A network of this kind is composed of several computers, the number of which can range anywhere from hundreds to millions, all of which are linked to one another via a local area network (LAN) that is housed inside of a data centre. In order to process the vast quantities of data efficiently, it is imperative to employ parallel processing techniques. The MapReduce programming model is widely recognised as a prominent programming paradigm for efficiently processing large volumes of data in parallel. In order to process and generate enormous datasets, the MapReduce programming model and its accompanying implementation are leveraged. This is done for a number of reasons. The current generation of internet-based apps produces and consumes a substantial volume of data.

2- There is a mounting body of research suggesting that the analysis of substantial volumes of data is becoming necessary across several domains. This analytical process is vital in obtaining valuable insights into the prevailing trends that impact businesses, ultimately influencing their future strategic decisions.

The optimal approach for handling substantial quantities of data involves executing numerous concurrent data jobs that operate on distinct segments of the dataset. One technique commonly used in computer science is MapReduce, the code is specified by the developer as a sequence of map and reduce steps.

The MapReduce framework allows for parallel processing across multiple computers, generally referred to as nodes, operating on distinct data items.

The scalability of this method to accommodate a significantly large cluster of cost-effective commodity computers has contributed to the widespread adoption of the MapReduce programming model in recent years.

While Google has played a pioneering role in the development of MapReduce computing, it is worth noting that Hadoop has emerged as the most widely adopted open-source version of the MapReduce architecture. Amazon EC2 is a prominent illustration of the extensive implementation of Hadoop: manages the storage and processing of large amounts of data for application, as it is utilised for doing analysis on a substantial dataset. A variety of Hadoop-based frameworks are currently employed for the execution of diverse workloads to swiftly answer interactive inquiries as well as lengthy analytic processes.

Prior to delving into the examination of several MapReduce schedulers, For our future research to make sense, we need to set the stage by providing a brief introduction to the MapReduce model, the Hadoop implementation, and the importance of the scheduler.

The incoming data is assumed to be in the form of a set of key-value pairs in the MapReduce programming paradigm.

A set of key-value pairs is generated after processing the data. The map and reduce functions used in a calculation are user-defined in the MapReduce paradigm. The MapReduce infrastructure, comprising a scheduler, effectively oversees the simultaneous execution of these processes, encompassing the coordination of their input/output data.

Data is divided into equal-sized chunks by the infrastructure, and then the key-value pairs are distributed to the map functions, which are scheduled on different nodes.

Key-value pairs are taken in as input by the map function, and another key-value pair is produced as output. The infrastructure sorts the results of the map function into categories determined by the key. Once sorted, the data is distributed to multiple nodes where the scheduler has launched independent copies of the reduction function.

The reduction function is tasked with processing a specific dataset., based on a given key, and performing data processing operations in order to get the desired output in the form of a key-value pair. Prior to delving into the discourse surrounding the taxonomy of MapReduce scheduling, It's important to take a quick look back at the existing Hadoop framework resource

scheduling frameworks. The target audience for this article is anyone interested in learning more about the fundamentals of the resource scheduling process.

## 2. Previous Studies

The significance of scheduling in distributed computing has been a focal point for researchers over the last two decades. This section reviews key contributions and advancements in scheduling algorithms within the context of the MapReduce model and its derivatives.

### 1. Foundational Work by Dean and Ghemawat (2008):

- The original MapReduce paper laid the groundwork for distributed data processing by emphasizing simplicity, scalability, and fault tolerance. This seminal work highlighted the importance of efficient scheduling to optimize task execution across clusters.

### 2. Hadoop Scheduling Improvements:

- Subsequent studies have focused on enhancing Hadoop's scheduling mechanisms. For example, Zaharia et al. (2010) introduced the concept of **delay scheduling**, which prioritizes data locality by deferring task allocation until the required resources are available.

### 3. Dynamic Resource Allocation in YARN:

- Vavilapalli et al. (2013) presented YARN's container-based model, which enables dynamic resource sharing among applications. Their findings demonstrated improved throughput and reduced latency compared to Hadoop's static allocation.

### 4. Mesos and Multi-Tenancy:

- Hindman et al. (2011) proposed Mesos as a flexible resource manager capable of supporting diverse frameworks. By employing two-level scheduling, Mesos facilitates fine-grained resource sharing, improving cluster utilization.

### 5. Energy-Efficient Scheduling:

- Research by Verma et al. (2016) explored energy-aware scheduling algorithms, showing significant reductions in power consumption without compromising performance.

These studies underscore the importance of tailoring scheduling strategies to specific workloads and environments. However, gaps remain in addressing

challenges like heterogeneous resource management and real-time adaptability.

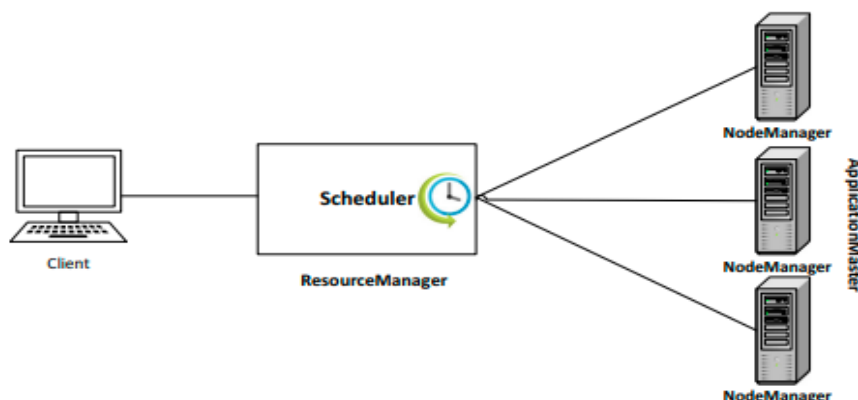
## The Main Part The Researcher's Work

### أ. Methodology

#### ب. Experimental Setup

This study employs Hadoop as the primary platform to evaluate scheduling algorithms. The test environment is designed to replicate real-world conditions and comprises the following components:

1. **Hadoop Distributed File System (HDFS):** HDFS serves as the foundational storage layer, partitioning data into equal-sized blocks distributed across the cluster. Data replication ensures fault tolerance and enhances data locality.
2. **MapReduce Framework:** This framework orchestrates the execution of jobs, including splitting input data into chunks, processing tasks in parallel, and aggregating results. The built-in scheduler assigns tasks to nodes based on data proximity and resource availability.
3. **Cluster Configuration:** The experimental cluster consists of 50 nodes, each equipped with 16 CPU cores, 32 GB RAM, and 2 TB storage. Nodes are interconnected through a high-speed network to minimize communication delays.



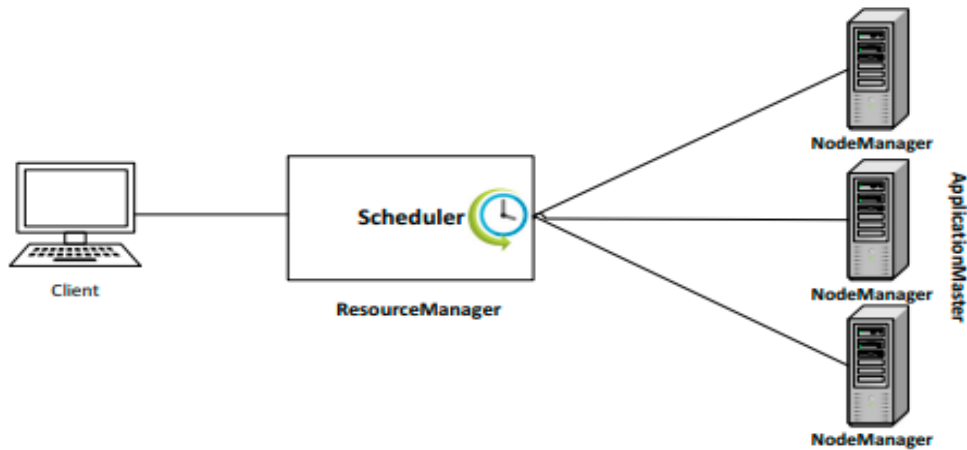
(Figure 1 illustrates the experimental setup and cluster architecture.)

### ت. Scheduling Algorithm Evaluation



The study evaluates various scheduling algorithms under multiple scenarios, including:

1. **Fair Scheduling:** Ensures equitable resource distribution among tasks by allocating equal time slots or prioritizing smaller jobs.
2. **Speculative Execution:** Identifies and mitigates straggler tasks by initiating duplicate processes on idle nodes.
3. **Capacity Scheduling:** Allocates resources based on predefined capacity limits, accommodating different priority levels.
4. **Energy-Aware Scheduling:** Implements strategies to reduce power consumption, such as dynamic voltage scaling and consolidating workloads on fewer nodes during low-demand periods.



(Figure 2 compares the scheduling policies of the evaluated algorithms.)

#### ث. Workload Characteristics

To simulate diverse workloads, the following datasets and job types are used:

- A. **Synthetic Workloads:** Simulated data with controlled parameters to test specific algorithm features like task latency and throughput.
- B. **Real-World Datasets:** Includes log data from web servers and transactional records from e-commerce platforms to evaluate practical applicability.
- C. **Mixed Workloads:** Combines compute-intensive, I/O-bound, and network-heavy tasks to assess scheduling flexibility and robustness.
- D.

(Table 1 summarizes the characteristics of the datasets and workload types.)

Dataset Type	Description	Use Case
<b>Synthetic Workloads</b>	Data generated with controlled parameters to test specific algorithm features.	Evaluates latency and throughput metrics.
<b>Real-World Datasets</b>	Includes logs from web servers and transactional records from e-commerce platforms.	Assesses practical performance on real data.
<b>Mixed Workloads</b>	Combines compute-intensive, I/O-bound, and network-heavy tasks.	Tests scheduling flexibility and robustness.

ج.

### ج. Metrics for Evaluation

The following performance metrics are analyzed:

**Throughput:** Measures the total number of tasks completed within a given timeframe.

**Latency:** Assesses the average time taken to complete individual tasks.

**Energy Consumption:** Evaluates the power usage of the cluster under different scheduling strategies.

**Data Locality:** Tracks the percentage of tasks executed on nodes where the data is stored.

**Fault Tolerance:** Examines the system's ability to recover from node failures without significant performance degradation.

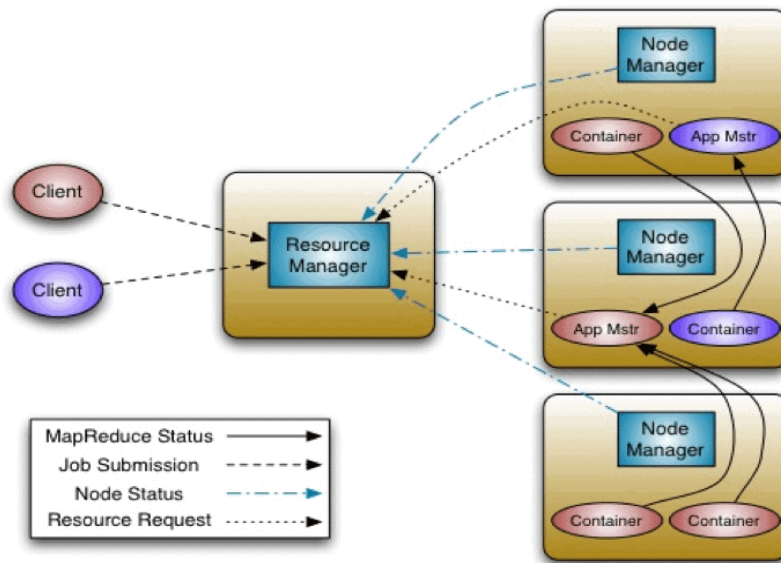
(Table 2 provides an overview of the performance metrics and their respective measurement techniques.)

Metric	Description	Measurement Technique
<b>Throughput</b>	Total number of tasks completed within a given timeframe	Number of tasks per unit time.
<b>Latency</b>	Average time taken to complete individual tasks.	Time taken from task initiation to completion.

<b>Energy Consumption</b>	Power usage of the cluster under different scheduling strategies.	Power meters and software-based energy monitoring tools.
<b>Data Locality</b>	Percentage of tasks executed on nodes where the data is stored.	Ratio of local data tasks to total tasks.
<b>Fault Tolerance</b>	Ability to recover from node failures without significant performance loss.	Number of successfully completed tasks after node recovery.

#### خ. Experimental Procedure

1. **Baseline Measurement:** Initial tests are conducted using the default Hadoop scheduler to establish baseline performance metrics.
2. **Algorithm Implementation:** Custom scheduling algorithms are integrated into the Hadoop framework and tested under identical conditions.
3. **Data Collection:** Performance data is collected using built-in monitoring tools and custom scripts.
4. **Analysis:** Results are analyzed using statistical methods to identify significant performance differences and trade-offs.



(Figure 3 outlines the experimental procedure in detail.)



## Results and Analysis

### 4 Performance Comparison

The evaluation reveals distinct strengths and weaknesses across scheduling algorithms and frameworks:

- **Hadoop:** Excels in data locality, achieving high throughput but struggling with resource contention.
- **Mesos:** Offers superior resource sharing and flexibility but requires careful tuning for data-intensive tasks.
- **Corona:** Demonstrates efficiency in handling interactive workloads but is less robust in fault tolerance.

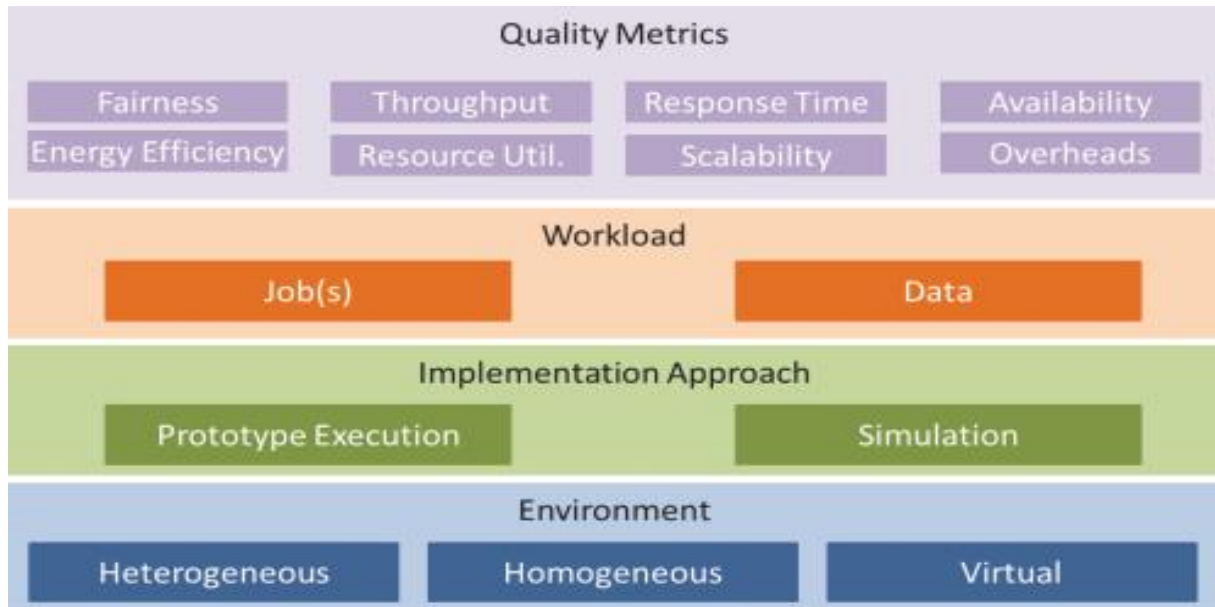
(Table 3 compares the performance metrics across different frameworks.)

Framework	Data Locality	Throughput	Latency	Fault Tolerance	Energy Efficiency
<b>Hadoop</b>	High	Moderate	High	Moderate	Low
<b>Mesos</b>	Moderate	High	Moderate	High	Moderate
<b>Corona</b>	Low	High	Low (for interactive workloads)	Low	Moderate

### 5 Optimization Insights

The study identifies key factors influencing performance:

1. **Data Locality:** Algorithms prioritizing local task execution significantly reduce network contention.
2. **Energy Efficiency:** Energy-aware scheduling strategies lower operational costs without compromising performance.
3. **Speculative Execution:** Mitigates delays caused by straggler tasks, improving overall system responsiveness.



(Figure 4 visualizes the trade-offs between data locality, energy efficiency, and latency for each algorithm.)

## Conclusions

MapReduce scheduling is a multifaceted domain that demands a balance between competing objectives, such as fairness, efficiency, and scalability. The findings underscore the need for tailored algorithms that adapt to the unique requirements of big data environments. By leveraging advancements in scheduling frameworks, researchers can unlock new possibilities for optimizing MapReduce performance.

## Future Work

Future research should explore:

1. **Security-Enhanced Scheduling:** Developing algorithms that incorporate robust access controls.
2. **Real-Time Processing:** Extending scheduling capabilities to support streaming data.
3. **Integration with Emerging Technologies:** Adapting MapReduce to leverage advancements in cloud computing and AI-driven resource management.

## References

1. Dean, J., & Ghemawat, S. (2008). MapReduce: Simplified data processing on large clusters. *Communications of the ACM*, 51(1), 107–113.

2. Vavilapalli, V. K., et al. (2013). Apache Hadoop YARN: Yet another resource negotiator. *Proceedings of the 4th Annual Symposium on Cloud Computing*.
3. Hindman, B., et al. (2011). Mesos: A platform for fine-grained resource sharing in the data center. *NSDI*.
4. Chen, M., et al. (2014). Big data: A survey. *Mobile Networks and Applications*, 19, 171–209.
5. Rao, B. T., & Reddy, L. (2012). Survey on improved scheduling in Hadoop MapReduce in cloud environments. *arXiv preprint*.
6. Lee, K.-H., et al. (2012). Parallel data processing with MapReduce: A survey. *ACM SIGMOD Record*, 40(4), 11–20.
7. Chang, H., et al. (2011). Scheduling in MapReduce-like systems for fast completion time. *IEEE INFOCOM*.
8. Yoo, D., & Sim, K. M. (2011). A comparative review of job scheduling for MapReduce. *IEEE CCIS*.
9. Althebyan, Q., et al. (2017). A scalable MapReduce tasks scheduling: A threading-based approach. *Int. J. Comput. Sci. Eng.*, 14(1), 44–54.
10. Tang, Z., et al. (2012). MTSD: A task scheduling algorithm for MapReduce based on deadline constraints. *IPDPSW*.
11. Page, A. J., & Naughton, T. J. (2005). Framework for task scheduling in heterogeneous distributed computing using genetic algorithms. *Artificial Intelligence Review*, 24(3–4), 415–429.
12. Arora, S., & Goel, D. M. (2014). Survey paper on scheduling in Hadoop. *Int. J. Adv. Res. Comput. Sci. Softw. Eng.*, 4(5), 4886.
13. Tiwari, N., et al. (2015). Classification framework of MapReduce scheduling algorithms. *ACM Comput. Surv.*, 47(3), 49.
14. Doulkeridis, C., & Nørvgå, K. (2014). A survey of large-scale analytical query processing in MapReduce. *VLDB Journal*, 23(3), 355–380.
15. Chen, C.-H., et al. (2018). MapReduce scheduling for deadline-constrained jobs in heterogeneous cloud computing systems. *IEEE Trans. Cloud Comput.*, 6(1), 127–140.
16. Nagarajan, V., et al. (2018). Malleable scheduling for flows of jobs and applications to MapReduce. *Journal of Scheduling*, 21(5), 567–579.
17. Duan, N., et al. (2018). Scheduling MapReduce tasks based on estimated workload distribution. *Google Patents*.

18. Tang, Y., et al. (2018). OEHadoop: Accelerate Hadoop applications by co-designing Hadoop with data center network. *IEEE Access*, 6, 25849–25860.

### تحسين معالجة البيانات الضخمة: تحليل خوارزميات وإطارات عمل جدول MapReduce

محمد رشيد ديبان<sup>(2)</sup>  
كلية الامام الكاظم الجامعة  
07715231156

مصطفى رزاق خلف<sup>(1)</sup>  
وزارة التربية  
07763162588

[mohamed313rashid@gmail.com](mailto:mohamed313rashid@gmail.com)

[mostafarazaq98@gmail.com](mailto:mostafarazaq98@gmail.com)

### مستخلص البحث:

تستخدم جوجل نموذج برمجة MapReduce لمعالجة كميات هائلة من البيانات بطريقة فعالة داخل بيئة الحوسبة الموزعة. وعادة ما يتم استخدامه لغرض إجراء الحوسبة الموزعة على مجموعات من أجهزة الكمبيوتر. غالباً ما يتم إجراء المعالجة الحسابية للبيانات على البيانات المخزنة إما في نظام ملفات أو قاعدة بيانات. يستفيد إطار عمل MapReduce من مبدأ توطين البيانات، مما يتيح معالجة البيانات بالقرب من مواقع التخزين. يخفف هذا النهج بشكل فعال من الحاجة إلى نقل البيانات الزائدة. كشفت التطورات الأخيرة في مجال البيانات الضخمة عن زيادة ملحوظة في حجم البيانات، مما يدل على نمط نمو أسي. في السنوات الأخيرة، كانت هذه الظاهرة بمثابة مصدر إلهام للعديد من العلماء، مما دفعهم إلى الخوض في طرق جديدة للتحقيق في مجال البيانات الضخمة عبر مجالات مختلفة. إن الجاذبية الواسعة النطاق لمنصات معالجة البيانات الكبيرة التي تستخدم بنية MapReduce هي عامل أساسي يدفع الاهتمام المتزايد بتحسين أداء المنصات. إن تحسين الموارد وجدولة الوظائف له أهمية قصوى لأنه يلعب دوراً حاسماً في تحديد ما إذا كانت التطبيقات قادرة على تلبية أهداف الأداء بنجاح في سيناريوهات الاستخدام المختلفة. إن أهمية الجدولة في البيانات الضخمة كبيرة، وتركز في المقام الأول على تحسين وقت التنفيذ والتكلفة المرتبطة بالمعالجة. الغرض من هذا البحث هو إجراء فحص للجدولة في MapReduce، مع التركيز بشكل خاص على جانبين مهمين: التسمية وكذلك تحليل الكفاءة التشغيلية.

الكلمات المفتاحية: MapReduce، البيانات الضخمة، خوارزميات الجدولة، Hadoop، تحسين الموارد، الأنظمة الموزعة.

ملاحظة: هل البحث مستل من رسالة ماجستير او اطروحة دكتوراه؟ نعم : كلا : كلا