

Replication-Based Query Management for Resource Allocation Using Hadoop and MapReduce over Big Data

Ankit Kumar, Neeraj Varshney, Surbhi Bhatiya, and Kamred Udham Singh*

Abstract: We live in an age where everything around us is being created. Data generation rates are so scary, creating pressure to implement costly and straightforward data storage and recovery processes. MapReduce model functionality is used for creating a cluster parallel, distributed algorithm, and large datasets. The MapReduce strategy from Hadoop helps develop a community of non-commercial use to offer a new algorithm for resolving such problems for commercial applications as expected from this working algorithm with insights as a result of disproportionate or discriminatory Hadoop cluster results. Expected results are obtained in the work and the exam conducted under this job; many of them are scheduled to set schedules, match matrices' data positions, clustering before determining to click, and accurate mapping and internal reliability to be closed together to avoid running and execution times. Mapper output and proponents have been implemented, and the map has been used to reduce the function. The execution input key/value pair and output key/value pair have been set. This paper focuses on evaluating this technique for the efficient retrieval of large volumes of data. The technique allows for capabilities to inform a massive database of information, from storage and indexing techniques to the distribution of queries, scalability, and performance in heterogeneous environments. The results show that the proposed work reduces the data processing time by 30%.

Key words: big data; hadoop; mapreduce; resource allocation; query management

1 Introduction

The global population exceeds 7.2 billion, and 2 billion people have Internet connectivity. According to McKinsey, 5 billion people use different mobile devices. With the revolution of this technology, millions

of people are generating huge amounts of information through the increased use of such devices. In particular, remote sensors consistently provide a huge volume of information that is neither constituted nor unmatched. Such information is known as big data^[1].

Three aspects that characterize big data:

- (1) The data cannot be categorized into regular relational databases.
- (2) The data are numerous.
- (3) The data are generated, captured, and processed very quickly.

Big data is a growing sector of the IT industry that focuses on business applications. It has sparked widespread interest in a number of fields.

Machine construction can be used in medical treatment, financial transactions, social media, and satellite imagery^[2]. Traditionally, the largest amount of data is possibly used to store structural information.

-
- Ankit Kumar and Neeraj Varshney are with the Department of Computer Engineering and Application, GLA University, Mathura 281406, India. E-mail: iiita.ankit@gmail.com; neeraj.varshney@gla.ac.in.
 - Surbhi Bhatiya is with the College of Computer Sciences and Information Technology, King Faisal University, Hofuf 31982, Saudi Arabia. E-mail: sbhatia@kfu.edu.sa.
 - Kamred Udham Singh is with the School of Computing, Graphic Era Hill University, Dehradun 248002, India. E-mail: kamredudhamsingh@gmail.com.

* To whom correspondence should be addressed.

Manuscript received: 2022-04-26; revised: 2022-07-07;
accepted: 2022-07-14

However, unstructured and semi-structured data are now driving data quantities. Because relational systems are used to manage databases for analyses, the translation of data stored in an unorganized and undesirable format may have an impact on intermediate processing.

Rapid growth of data rate:

YouTube:

- (1) Users upload 100 h of fresh videos each minute.
- (2) Monthly, over 1 billion particular users access YouTube.
- (3) Six million hours of video have been observed each calendar month, equating to nearly an hour.

Facebook:

- (1) Every moment, 34 722 likes are enrolled.
- (2) One hundred terabytes (TB) of information are uploaded every day.
- (3) The website has 1.4 billion consumers. Moreover, the site has been translated into 70 languages.

Twitter:

- (1) The site has more than 645 million consumers.
- (2) The website creates 175 million tweets daily.

Google:

The website gets more than 2 million search queries per second daily, and 25 petabytes (PB) are refined.

Four-square:

- (1) This website is used by 4–5 million individuals globally.
- (2) This website gets more than 5 billion check-ins daily.
- (3) Every moment, 571 new sites are found.

Apple:

Approximately 47 000 programs are downloaded each second.

Tumblr:

Weblog owners release 27 000 new articles per second.

Brands:

More than 34 000 likes are enrolled per second.

Flickr:

Users upload 3125 brand-new photographs per moment.

Instagram:

Users discuss 40 million photographs daily.

WordPress:

Close to 350 sites are published each second.

LinkedIn:

2.1 million collections have been made.

In 2020, 73.5 zettabytes were created, it was 2.52 zettabytes in 2010, data on worldwide output have increased by 4300% by 2022. Data exploding over

the Internet today & developers and Internet users can easily discuss the importance of data for their project and applications. The Internet has become an enormous place where one can easily find everything from a needle to human genome reports. The Internet is a complex mesh of tools, frameworks, applications, algorithms, and hardware commodities for storing, processing, and managing data worldwide. A single data server or application is adequate for handling data generated by a single user. Choosing which applications to use or how to manage data is an ongoing research project. Today, the Internet is inextricably linked to data, but the more general term for this type of massive data would be big data. In this scenario, another question is how such data are generated^[3,4]. Big data frameworks, applications, algorithms, and hardware commodities allow people all over the world to store, process, and manage data.

1.1 Hadoop ecosystem

Hadoop was built to scale up computers that perform processing from single to thousands of nodes, where each node or machine can work as a local computation or storage drive. The framework of Hadoop has four significant modules, as depicted below, which divide the functionality-based working of Hadoop into distinct units. The four modules that are most important in Hadoop^[5] are as follows:

- Hadoop system (standard)
- Hadoop MapReduce
- Hadoop Distributed File System (HDFS)
- Hadoop YARN

Other common softwares, such as Hive, Pig, Scala, Scoop, and CDH, facilitate data input/output (I/O) or querying into Hadoop. MapReduce is used to process data on Hadoop or HDFS, but most of the functioning is related to coding in Java, Ruby, Perl, or Python. Most developers prefer querying languages like SQL to talk to a system. Hive or Pig, both open-source implementations, is used to facilitate the interaction between Hadoop and developers knowing querying languages shown in Fig. 1^[6].

The Hive interpreter converts a query into a MapReduce code and interacts with the HDFS. The Pig interpreter turns it into a simple scripting language that again communicates with MapReduce and then to the HDFS. Pig and Hive can work very well, but the overall gist of this system is that it takes a lot of time to convert queries into MapReduce forms and then interact with Hadoop. It could incur substantial cost losses when

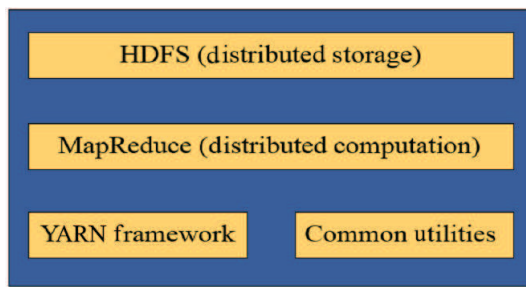


Fig. 1 Hadoop framework architecture.

performed for significant amounts of data. Another creation was Impala, designed to interact directly with HDFS by converting SQL queries into HDFS format. Impala is optimized for low-latency queries, a crucial requirement in Hadoop and big data scenarios. Typically, an Impala query runs way faster than a Hive or Pig query because of its low-latency structure and direct interaction with the HDFS^[7].

1.2 Homogeneous data clusters in Hadoop

The clustering in Hadoop systems is essential for planned research works. Similar clusters in a group make a collection of the same type of machine for order processing and warehouse data collection in the HDFS. The hardware similarity in homogeneous clusters makes the job processing and data storage mechanisms easy to implement and deploy. Hadoop and big data^[8] work well in a similar cluster environment.

1.3 Working in Hadoop

Hadoop has been trying to get users' and developers' praise. Its framework allows users to rapidly write and test a distributed system. Hadoop's MapReduce framework is capable of the reliable and fault-tolerant mode of the application, which can handle a significant amount of data in large parallel clusters of commodity hardware writing. Hadoop MapReduce is a program that can perform two basic tasks: map task and reduce task.

1.4 MapReduce

Serially processing a job takes longer than necessary, which can cause delays in other system functions, such as I/O, querying, or data retrieval. As a result, MapReduce is designed to process data in parallel. The MapReduce framework is in charge of task scheduling, monitoring, and re-running failed tasks^[9]. It has a master-slave hierarchy. Each node is labeled "JobTracker" master and "TaskTracker" slave. Each node is comprised of a single JobTracker and TaskTracker, with JobTracker being primarily responsible for resource management, job task

slave planning, and resource availability/consumption tracking. JobTracker also tracks job components, tasks, monitors, and re-executes failed tasks. Because they regularly task status information, JobTrackers also perform the task assigned to them, following the TaskTracker design. Mappers and Reducers work on the processing, as explained below. A mapper node receives a large chunk of data assigned to a reducer.

2 Literature Review

In this section, we review the state-of-the-art big data analysis methods.

2.1 Big data processing

Malik et al.^[10] revealed the crucial importance of big data in today's world. As far as applications are concerned, big data have a lot of usages but pose challenges. They have used the concept of 4 Vs to describe big data, including veracity. They have addressed some fundamental problems that need to be catered to so that problems can be converted into opportunities. They have reported advances, research on big data, and six significant challenges encountered during the processing of unstructured data. They have also proposed the research problems still open in the big data community.

In addition to the scientific issues discussed in detail, the authors have mentioned open problems that arise in the big data scenario. The high dimensionality of big data, subsampling of data in Hadoop using MapReduce, computational complexity of a dataset, distribution issues in real computing, problem of processing non-structured data, and visualization of datasets to improve efficiency are some of the open challenges that the scientific and development community is grappling with.

Cloud computing applications used today involve processing a significant amount of data. Such data could be in different formats, which require considerable processing. As mentioned by Yao et al.^[11], the data processing rate is exponentially growing, but the same is not right for computing power. The need for new strategies is much awaited, either an improvement of reliable storage methods that remain inexpensive or the development of new tools for analyzing structured and unstructured data. The focus is that different file systems have issues in data migration. On a heterogeneous cluster in Hadoop, the load-sharing task is executed to improve processing. The jobs are processed with a fast execution time on a high-speed end node, but the data are processed

much more slowly on a low-speed end node. The load-sharing mechanism enables the high-speed end node to relieve the low-speed end node, but when the volume of data to be processed is high, the load-sharing mechanism is bound to fail. The overhead of the data movements of unprocessed data between the slow-speed end node and high-speed end node affects the performance of Hadoop in different clusters.

Alshammari et al.^[12] presented another critical issue of heterogeneous Hadoop clusters, which is the namespace limitations offered by the systems. The namespace storage is located at the RAM, and as the data are stored in file system metadata formats, the namespace assumes names similar to the data nodes. The problem arises with sharing name nodes and maintaining metadata over the name nodes in a Hadoop environment for a vast number of files. An individual namespace server has its limitations in handling and storing data. The surveys conducted reveal that only 200 bytes are needed for a name node server to a single object file. The fact taken into consideration shows that when this is multiplied with other data blocks residing on the same server, we obtain a vast number of metadata to store. This is one of the critical attributing factors to the performance of name node servers in Hadoop clusters. The name node needs RAM to process requested queries and fetch I/O files, among others. When the space is entirely consumed by the name node's metadata only, one can easily understand the overheads of the processing namespace in name nodes. The whole cluster becomes unresponsive when the name node is busy processing internal job requests. Data-intensive job applications must compromise and compete for resource distribution on name nodes. The data transfer requested by client nodes will have to suffer because each request is first processed on the name node, where internal metadata are managed and asked by the client node. These issues are brought to light by the authors, who convey a great insight into how a simple name node and cluster configuration can affect the proceedings of Hadoop in cloud computing.

2.2 Hadoop and MapReduce

Ullah et al.^[13] raised concerns about the challenges that are posed by supercomputing communities. They brought forward issues that arise when data are encountered in petabytes or zettabytes. The vast amount of data created because of numerical simulations as the resultant of supercomputing operations is the key

area to focus on. The parallelism in big datasets could be viewed from three different aspects, namely, data applications that involve a high rate of inbound and outbound movements, server applications that involve a network with high bandwidth to process all queries, and scientific computing applications that need high speed of processing and memory performances to decode complex phenomena. Data parallelism, if appropriately achieved, could benefit any organization at a significant level. The current trends in the IT industry have proven its high profitability because of the data industry.

In the research by Parmar et al.^[14], virtualization is an excellent way to harness the capabilities of Hadoop and MapReduce on data clusters^[8]. Data-intensive jobs on nodes could be significantly improved by incorporating virtualization as a green IT factor. Resource consumption could be reduced and save a lot of energy. The work describes the challenges met in the virtualization environment for Hadoop and MapReduce tasks. It also describes schemes to overcome and manage it, so virtualization is proven as a practical solution for the big data crisis. The authors also proposed a comprehensive rating scheme that will check the performance of clusters in their respective environments. This scheme aims to extend YARN for configuring Hadoop and MapReduce. The experiments conducted under this niche reveal that performance improves through an integrated performance model. The research has a wide scope, which includes automating adjustments needed in Hadoop configurations. The experimental setup consists of multiple nodes, and different datasets were deployed on them to test the analysis. The nodes were divided into different cluster sizes, and TestDFSIO was used. The tasks depict a high rate of performance and parallelism. They have mapped the task on different process to test the reduction in similarity.

Kumar and Singh^[15] mentioned that the distinguished, profiled critical execution of queries has equitability issues. For the existing MapReduce-based information warehousing framework, prediction-based systems are suggested, and an inquiry planning framework was built, which scaffolds that a semantic hole exists between MapReduce runtimes. Furthermore, the inquiry compiler empowers productive inquiry planning for quick and reasonable enormous information analytics. The author performed experiments to overcome the drawbacks of design and improve the performance of large-scale clusters. A cross-layer scheduling framework was used

so that Hive queries can be percolated and semantically extracted, and a multivariate execution time prediction and two-level query scheduling were performed^[11].

2.3 Query performance

Huang et al.^[16] mentioned that a good query performance prediction is nearly impossible with old query performance models. For effective query optimization, resource management, and user experience management, the execution latency used by common analytical cost models is a poor predictor of comparing candidate plans' costs. QPP seems a practical approach for the high success of predictive tasks that can generate good results within new performance models. The work supports the idea that prediction can be studied from different perspectives, such as analyzing the granularity at operator levels. The proposed predictive modeling technique manifests the differences in the query execution behavior at varying granularity levels to understand the need and usage of coarse-grained or fine-grained operations.

The research performed by Soualhia et al.^[17] specifies the use of similar I/O techniques to manage scientific applications^[10]. They incorporated a data sieving and two-phase I/O technique to optimize the I/O throughput. The proposed work emphasizes offering parallel access to shared datasets for heavy workloads. They also used cross-layer optimizations to improve the I/O throughput for scientific applications, such as GEOS-5. Although the proposed work is implemented for scientific applications, it presents a broad idea of how datasets of dynamic nature behave. It also points out the significance of real-times change in queries, which is a significant issue in Hadoop, MapReduce, and big data scenarios. The I/O performance of computing can be easily categorized into different levels based on file systems and various needs.

Analyzing data that rapidly grows and changes has driven many new technologies and frameworks to involve predictions and new learning methods. The rootworker of these technologies is algorithms that help categorize, assess, and analyze essential data from meaningless information. For example, in an online store's log, the time spent on a product is useful information, whereas surfing random products is not that useful. To understand this vital key of details from PB of data, algorithms need to be smart and efficient. Here, clustering algorithms have become a powerful resource to manifest essential data from clutters. Tao et al.^[18]

proposed comparative studies of existing algorithms used for big data classification of datasets. The survey conducted in this work classifies the functionality of candidate algorithms primarily used in big data regarding theoretical and empirical analyses. The parameters taken for considering the gross efficiency of an algorithm involve many internal and external factors. They are performance metrics, runtime, scalability tests, and stability. They also described the best possible clustering algorithm used in big data.

Tao et al.^[18] supported the fact that manifesting digitized data can be a key to solving many problems posed by the Internet. Network traffic patterns and social media data, among others, could be beneficial for corporations and people. However, the processing of such large-sized data also comes with deficiencies. It also requires a significant amount of hardware and processing power. Going through all these data can be time-consuming and challenging to analyze. Clustering techniques are used to manage all these problems and provide solutions to efficiently store and analyze information. Clustering algorithms for big data have to consider the volume, velocity, and variety for a well-thought analysis. Therefore, the 3 Vs are the core of big data characteristics and the key to designing a clustering algorithm. The problem that arises in algorithm clustering is deciding what type of algorithm to use. The choice of algorithm for different datasets differs as every dataset has its ups and downs.

Qin et al.^[19] outrightly discussed their ideas on parallel data processing in heterogeneous clusters. They recommended that Hadoop is in dire need of mechanisms to support scalability and improvement in performance. They evaluated methods to consider the view of storage and techniques for indexing and other processes related to queries, such as distribution, performance, parallelization, and scalability, in a heterogeneous environment. The authors proposed designing Hadoop's architecture so that various data nodes can be scaled up to optimize Hadoop clusters. The test results reveal that the replication factor of data nodes can experience node failures in adverse scenarios. They claim that although small data jobs can be handled by other methods, MapReduce is superior when dealing with large datasets. To address the challenges posed by data-intensive jobs, the cost of hardware computing and mixing software was set. The authors clearly stated that the data size has increased in recent years from a few gigabytes to petabytes. This increase in data generation is due to the

increased use of content management systems, social media, logs, and other similar tools. This volume of data necessitates the use of more than a couple of machines to manage it for all Internet users. As a result, handling these generated data across heterogeneous clusters is critical for harnessing the power of Hadoop and the big data platform. The HDFS is an excellent environment for storing and maintaining data for a large number of users. The main issue is saving such data and keeping them secure. They must also be processed in acceptable time frames with efficiency and speed. The authors proposed a new design model to meet future system requirements and help organizations find the right partner for large data storage and processing.

2.4 Distributed processing and MapReduce

Wang et al.^[20] stated that MapReduce is essential for the distributed processing of large-scale applications. Hadoop is used in data-intensive jobs, such as web indexing and data mining because it completes tasks with slow response times. It makes them very widely available, but a problem with data placement is often encountered. Data placement on data nodes becomes a disarray job in heterogeneous Hadoop clusters. Hence, there must be a good way to assign data to nodes so that high-end nodes get jobs that require fast processing and slow-end nodes get appropriate time and jobs to process based on their capabilities. The authors devised methods for implementing data placement in heterogeneous clusters so that nodes are viable enough based on their capacity. Virtualized data centers face challenges because data locality is not taken into account for speculative map tasks. Because most maps are assumed to be stored locally, they are insufficient for virtualized centers. By specifying a design constraint on data nodes for data placement in Hadoop's heterogeneous clusters, the performance is greatly improved. The performance of data-intensive jobs could be greatly improved using various data placement techniques. Rebalancing data nodes as needed to perform the functions of a data-intensive job can improve performance, reduce latency, and provide a stable system. Heterogeneous Hadoop clusters can be used for data-intensive jobs by managing data placement locally and smartly adapting to the environment so that MapReduce jobs can be executed quickly and efficiently.

MapReduce algorithms have a wide range of applications, as mentioned repeatedly. They have emerged and continued to do so for the open-source

implementations of Hadoop for millions of users handled at once^[21]. The performance of Hadoop clusters is closely related to the short response times as processed by MapReduce^[11]. In homogeneous clusters of Hadoop, tasks take a linear progression model, but this case is not the same for heterogeneous clusters. The authors proposed a new algorithm, i.e., Longest Approximate Time to End, to improve the performance degradation of heterogeneous clusters due to impromptu task scheduling. The algorithm in the work has improved the performance twice on a large cluster. The response times of heterogeneous clusters can be enhanced in such a manner by proposing a scheduler algorithm to identify tasks and manage them accordingly on data nodes. The benefits of Hadoop and MapReduce are that programmers need not worry about faults occurring at the backend. MapReduce works all right if a node crashes but does not let the task run in disarray. The proposed scheduler works to maximize performance and improve robustness.

2.5 Classification approaches

In the research by Li et al.^[22], the classification of electronic documents was used to analyze applications. The type of classification used here is naïve Bayes because it is simple, efficient, and effective for significant data documents. Compared to some statistical methods, it has certain limitations, but it can help categorize long documents into more straight sets. The authors proposed a parallelized semi-lazy naïve Bayes approach to speed up the processing to a very high rate. Automatic Document Classification (ADC) is the key component used by the authors to experimentally conduct real case scenarios that help understand the nature and demeanor of datasets.

The ADC methods incorporate artificial intelligence in their initial stages, where the set of documents is used to learn and train. Once a sufficient data learning size is reached, they are used to design techniques to categorize related neighbors. They presented a good insight into how data should be classified by first learning from datasets and then applying the same to those. They somewhat resemble nearest-neighbor procedures and, in a sense, Bayesian models. The method decomposes the clustering problem into a hierarchical model where the hierarchy is divided into parent and child. However, there is a super parent and favorite child who are dependent on each other. A validation set is used to decide the heuristics and find the correct prediction for

measuring the details. The metric selection depends on the validation set, which enables the building of the network and deciding on the super parent and favorite child set. This method, commonly called “SP-TAN”, was popular in small datasets. The datasets used in ADC are not relevant when used in combination with SP-TAN. The dependencies in a large dataset are cumbersome to manage, and the identification process becomes tedious. When calculated, the computational costs determine that the performance cannot be overloaded and contented in the case of TAN or SP-TAN functioning, leading to a significant increase in the computation time and thus losing performance in the category of large datasets.

Wang and Cao^[23] also supported the use of naïve Bayes in their clinical experiments to understand the neuroimaging framework of humans. The research involves different stimulus phases, but the data collection and classification presented a highly intelligent insight into how data classification can be managed if performed correctly. Every response gathered as a stimulus provides a relatable experimental platform to determine the right heuristic technique for the data classification of straightforward tasks with complex environments.

2.6 Text clustering

Qureshi et al.^[24] followed the model framework of Hadoop’s MapReduce and calculations based on the Term Frequency-Inverse Document Frequency (TF-IDF) and its cosine similarity. The text clustering on which the method is applied is usable in many applications that are data-intensive. They are first collected and TF-IDF is applied. Then, a cosine similarity matrix is used to cluster them into available components. The method delivers efficient and faster results compared to other Hadoop’s MapReduce framework algorithms.

Clustering done pairwise has many benefits, as pointed out by the authors. It can be used to cluster documents/text quickly and arrange subsequent processing into easy takeover turns. The phases in the Cosine Similarity with MapReduce algorithm include building a vector model of data, then calculating the TF-IDF, and producing the angle of cosine for similarity in matrices of pairwise texts.

3 Proposed Algorithm

This section contains the proposed work and model and discusses the parameters used to compare the results.

3.1 MapReduce programming model

To reduce the programming model during mapping, the

accounting method is divided into two main stages, namely, the duration of the mapping process and the level of rearrangement. After the size expires, the level of rearrangement begins. Map efficiency is reduced using a reducer action and mapper activities. However, a mapper task will be split into an input signal, and the map pointers will be executed in parallel with pointers and carry the map’s properties. Measurement activities will create several effects. Before the reducer can perform work, the machine automatically populates the sequence of the population and intercepts it with the concealed reducers^[25]. Then, a reducer task will generate output files that create an output signal, perform a decreasing function, and be merged for results by reducing the event. Programmers need to compile the purpose of a mapper and reducer:

Map: (input text) \rightarrow {(key i , value i) | $i = 1, \dots, n$ };

Reduce: (key, [value 1, . . . , value n]) \rightarrow (key, final value).

The task failure and node are likely to occur within a MapReduce job implementation procedure. All mapper activities will proceed on available nodes whenever a node fails. This type of system, i.e., rescheduling, is easy but introduces a good deal of time price for customers with responsiveness prerequisites, and its effect is not acceptable.

3.2 Improved rescheduling algorithm

In order to make changes to the frame, a different rescheduling technique is used, which helps to re-establish the system so that it has a high level of performance. The activity and node failure is taken into account, and delays are reduced. Two types of documents, i.e., checkpoint files and index files, are introduced, which are international documents created prior to implementing a single-sized task^[26]. The sanctuary is in charge of preventing the execution of the current work of checkpoint documents. Whether or not the job failed, it can be continued through the service 1 node using the information in the checkpoint file. In addition to index documents, the sanctuary is in charge of publishing jobs that help reduce current results. In the event of a failure, the index file can be used as a file. The algorithm is divided into two parts: the controller node’s worker node and other functions^[27]. Furthermore, because the controller node is critical, it is necessary to maintain a completely consistent “hip copy” of the two elements of the algorithm.

As soon as an activity failure does occur, the node reads the checkpoint file stored in the disk, re-establishes

the task status into the checkpoint, and then reloads the results generated until collapse. Therefore, re-execution is avoided.

After a failure occurs, the scheduler must schedule the job with the master node and disrupt mapper activities to replicate nodes, as discussed in Algorithm 1. It may create the consequences of activities using the global indicator document to reduce the period. When nodes and tasks collapse until the checkpoint, the advancement will continue from the checkpoint. Of course, if the act of rescuing the checkpoint is neglected, the advancement will soon begin with the checkpoint. The simple failover plan is cost-effective in a distributed computing environment. The frequency of the rescuing checkpoint must be chosen: while there will be a frequency at the case, a high frequency will most likely offer the cost for great cost savings and failover for running, as discussed in Algorithm 2. In our experiment, the frequency value is assigned to a single checkpoint for every 105 pairs after fixing^[29].

The actions onto the node are rescheduled to replicate habitats in case a collapse occurs in the reducer point. If mapper activities are finished and the consequences have been duplicated to the replicate node, there is not any requirement to replicate the mapper activities over the node upon which the completion time of the MapReduce

Algorithm 1 Controller node to improve the performance of the Hadoop framework

Input: Number of Job 1, 2, ..., n

Output: Master node job schedule

Step 1: The controller node predicts that the measurement is performed with reducer actions in different worker documents.

Step 2: Choose the replicas for each job node.

Step 3: Check out the results of all staff documents.

When all the results are available, these results combine the results and indicate that the project is completed.

Otherwise, stay in Step 3 and wait.

Step 4: Slowly, all activation nodes send search packs.

1. If most staff respond NO, explore from Step 4 to Step 1 and then stay still.
 2. If a node does not react in a given time interval, subsequently inform the node as indicated.
 3. Employee ID and node will get almost incomplete tasks.
 4. Establish all the perfect map operations on the global line and reschedule their renewal nodes.
 5. You can determine the neglected activities of the failed node. After rectifying it, the group asks for re-established nodes that have intermediate facilities without renewing the monitor's functions.
 6. Once all nodes are completed, insert additional unexpected activities into the node.
-

Algorithm 2 Slave node to process the data

Input: Different sizes of Job 1, 2, ..., n

Output: Job execution time

Step 1: Assess the kind of specified failed task.

(1) If it is indeed a sizable work, it assesses a new attempt or all the failed work to evaluate what is evaluated.

(a) If it is a new initiative, launch it, and implement it.

(b) If the area fails to rehabilitate, buy and implement its progress from the nearby checkpoint document.

(c) Suppose that it can actually be redistributed from the dirty failures and nicks other than the failed works. In this case, the global index document will examine the failure of the work and immediately rearrange the intermediate impact using the information inside the global index document.

(2) If it is indeed a reducer task, it is a new commitment, i.e., whether this neglect removes the failure of the work or determines the incredible work from various nodes.

(a) If it is a new initiative, it is initially implemented and implemented.

(b) If this only reinstates the neglected works from different noses, browse the intermediate information from the local disk and then implement it.

(c) If it is performed unexpectedly from different nodes in a brand-new number, it will browse through the information provided in the document and implement it.

Step 2: Produce a localized checkpoint file and worldwide index document for any mapper task.

Step 3: Begin the mapper task.

(1) After the node's memory processing of the map, the intermediate information among the locals is empty. After the sink ends, place the input and map's location (position, then Map 1, where Map 1 is the key pair value) area checkpoint file^[28].

(2) Add a key-value pair corresponding to the positioning supply, which produces a key-quality pair, and two different strategies are distributed among the global index documents.

(a) Enter the key-quality pair output and list it permanently (T1, offset) in the global index document. Hence, only the offsets will need to be processed during re-establishment. T1 usually means that it is an inch album type.

(b) To insert a pair that indicates a result, list the opportunity as (Offset 1, Offset 2), so that input signals associated with Offset 1 and Offset 2 do not have an output signal and will be run again. T2 means that the fact is among the two sorted facts, where T2 represents the Task 2.

Step 4: After a mapper action ends, repeat and deliver the intermediate leads on the reducer nodes. Copy the information needed from the reducer's completed job to replicate nodes, inform the conclusion of the mapper endeavor to the perfect node, and then delete the indicator document and checkpoint file.

task is significantly diminished.

3.3 Algorithm for query optimization

This algorithm stores the query in a query buffer and

breaks the query issue by the job scheduler in the Hadoop process. First, we sort all queries in the query buffer based on keywords. The user enters the keyword-based query condition in the keywords field. In addition, the user has the following options:

- (1) Choice of “AND” or “OR” Boolean stringing of keywords;
- (2) Specification of the number of resulting objects to be shown on each page;
- (3) Selection of partial matching or exact matching of keywords;
- (4) Choice of how much information is desired to be presented in the result.

The access path for a simple query is generated by utilizing information brokers who support keyword-based searching.

Query processing is currently composed of a five-step process:

Step 1: Source selection;

Step 2: For each source, the translation of the simple query forms data into a valid URL address to the source;

Step 3: Parallel execution of each subquery against an individual source, translation of each subquery result into the consumer’s preferred result representation, and merging of subquery results from different sources into the final result format;

Step 4: The query management results are produced by combining the search results from all the available data sources selected to answer the query issue by the job scheduler process.

4 Result and Discussion

The following results have been established by directing the proposed methodology in the system. The heterogeneous Hadoop clusters were used to test the working of big datasets, which emulated the I/O throughput, the average rate of execution, standard deviations, and test execution time^[30].

4.1 Comparison of productivity among Hadoop schedulers

A comparison among the most efficient schedulers used in Hadoop, i.e., FIFO, HFS, and HCS, was performed. The following graphs establish that schedulers can overtake other schedulers used in the HDFS and internal mechanisms to share a file or job over several default systems to execute and store them reliably.

Figure 2 shows a trade-off among various schedules, which may lead to the optimized performance of systems in the HDFS for clustering big data shown in Table 1, and the complete process is discussed in Algorithm 3.

The response time was affected by data loads, which were the logs of an e-commerce website and response time taken by schedulers in mapping and reducing them to store them on the DFS and manually extracting and managing them. It also establishes that the amount of data is general in big datasets. Hadoop can optimize the processing instead of frequent data loading and unloading to obtain better insights about data, as shown in Fig. 3. This process also proves that once data are accepted by a node, the average time to consume and process them may vary depending on the type of scheduler being used^[31].

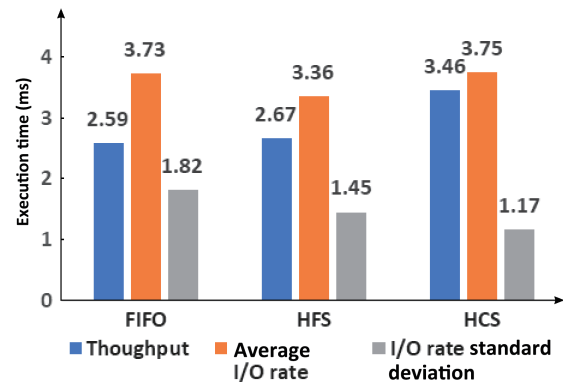


Fig. 2 Comparison of different Hadoop schedulers.

Table 1 Comparative study of exiting vs. proposed work.

| Scheduling algorithm | Data size | Previous methodology | Proposed methodology |
|----------------------|-----------|----------------------|----------------------|
| | | execution time (ms) | execution time (ms) |
| FIFO | 500 MB | 4000 | 3950 |
| | 1 GB | 7800 | 7333 |
| | 2 GB | 9900 | 9165 |
| | 5 GB | 14 500 | 11 130 |
| | 10 GB | – | 12 780 |
| HFS | 500 MB | 3950 | 3925 |
| | 1 GB | 7200 | 7065 |
| | 2 GB | 9200 | 9085 |
| | 5 GB | 13 500 | 10 980 |
| | 10 GB | – | 12 400 |
| HCS | 500 MB | 3900 | 3875 |
| | 1 GB | 7000 | 6998 |
| | 2 GB | 9000 | 8912 |
| | 5 GB | 12 000 | 10 260 |
| | 10 GB | – | 12 610 |

Algorithm 3 Query optimization

Input: Keyword query $Q = \{k_1, k_2, \dots, k_m\}$, list of related keywords $k_1^i, k_2^i, \dots, k_n^i, 1 \leq i \leq m$

Output: Q_1, Q_2, \dots, Q_k , where $Q_i = (q_1^i, q_2^i, \dots, q_m^i)$

Step 1: Delete redundant related keywords

$\{k'_1, k'_2, \dots, k'_g\} \leftarrow \{k_{11}, k_{12}, \dots, k_{1n}\} \cup \{k_{21}, k_{22}, \dots, k_{2n}\} \cup \dots \cup \{k_{m1}, k_{m2}, \dots, k_{mn}\} (g \leq m \times n)$.

Step 2: Build a Viterbi model $\lambda = (A, B, \pi)$

$A_{g,g} = \{a_{ij} \mid 1 \leq i, j \leq g\}, a_{ij} = \text{simi}(k'_i, k'_j),$
 $B_{g,m} = \{b_{ij} \mid 1 \leq i \leq g, 1 \leq j \leq m\}, b_{ij} = \text{simi}(k'_i, k'_j),$
 $\pi_i = (k'_i) = \text{simi}(k'_i, k_1)$.

Step 3: Initialize

//Variable $\delta_t(i)$ as the maximum probability in all paths whose state is i at time t .

//Variable $\psi_t(i)$ is the $(i - 1)$ -th node of the path with the maximum probability in state i at time t .

$\delta_1(i) = \pi_i b_i(k_1), i = 1, 2, \dots, g$

$\psi_1(i) = 0, i = 1, 2, \dots, g$

Loop

For $t = 2, 3, \dots, m$

$\delta_t(i) = \max\{1 \leq j \leq g\} [\delta_{t-1}(j) a_{ji}] b_i(k_t), i = 1, 2, \dots, g$

$\psi_t(i) = \arg \max\{1 \leq j \leq g\} [\delta_{t-1}(j) a_{ji}], i = 1, 2, \dots, g$

Return $(i) \cap \psi_t(i)$

End of loop

End of query buffer

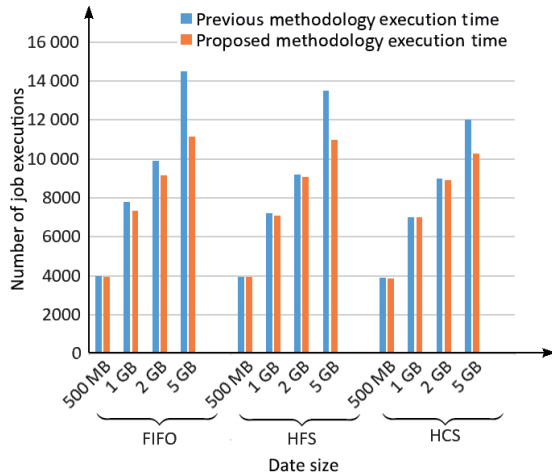


Fig. 3 Comparative performance analysis of previous methodology vs. proposed methodology.

4.2 Comparative analysis of the Hadoop framework with different schedulers

In this section, we have compared the performance of FIFO, HCS, and HFC in Tables 2, 3, and 4, respectively.

Figure 4 represents the comparative study of response time computation vs. data size using FIFO algorithm.

Figure 5 represents the comparative study of response time computation vs. different data sizes using HCS algorithm for query processing. Figure 6 depicts the

Table 2 Comparative study of FIFO job scheduler in Hadoop without query improvisation.

| Data size (MB) | Completion time (ms) |
|----------------|----------------------|
| 100 | 190 |
| 200 | 410 |
| 300 | 380 |
| 400 | 410 |
| 500 | 610 |
| 1500 | 2000 |

Note: Average performance ratio = 0.75.

Table 3 Comparative study of HCS in hadoop without query.

| Data size (MB) | Completion time (ms) |
|----------------|----------------------|
| 100 | 93 |
| 200 | 123 |
| 300 | 210 |
| 400 | 368 |
| 500 | 387 |
| 1500 | 1181 |

Note: Average performance ratio = 0.78.

Table 4 Comparative study of HCF in Hadoop without query.

| Data size (MB) | Completion time (ms) |
|----------------|----------------------|
| 100 | 88 |
| 200 | 188 |
| 300 | 218 |
| 400 | 320 |
| 500 | 410 |
| 1500 | 1301 |

Note: Average performance ratio = 0.86.




Fig. 4 Comparative study of FIFO job scheduler in Hadoop without queries.

Query processing time analysis of different data sizes using the HFC job scheduling algorithm.

4.3 Performance evaluation

Only performance test results were evaluated with clustered nodes and are shown in Table 5. The default

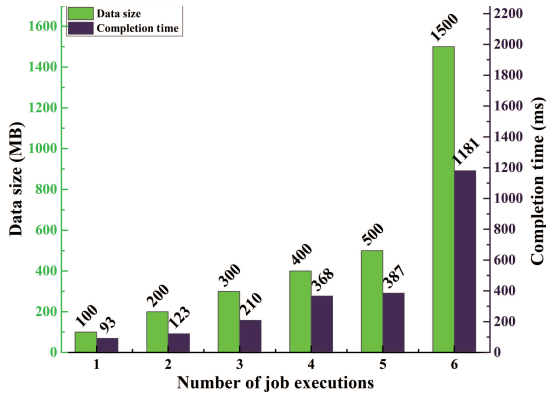


Fig. 5 Comparative study of HCS job scheduler in Hadoop without queries.




Fig. 6 Comparative study of HCF job scheduler in Hadoop without queries.

Table 5 Comparative analysis of Hadoop framework for different files’ sizes with their execution time.

| File size (MB) | Execution time of default Hadoop | Execution time of proposed method |
|----------------|----------------------------------|-----------------------------------|
| 12.2 | 00:00:41 | 00:00:40 |
| 36.5 | 00:00:53 | 00:00:53 |
| 146.0 | 00:01:51 | 00:01:52 |
| 300.0 | 00:02:09 | 00:01:53 |
| 500.0 | 00:02:30 | 00:02:07 |
| 938.0 | 00:04:42 | 00:04:07 |
| 1500.0 | 00:06:07 | 00:05:03 |
| 2000.0 | 00:08:51 | 00:07:17 |

option Hadoop and Hadoop contexts performed by the system were always implemented through the collection of information. The size of the data collection is 12 MB, 36 MB, 300 MB, 500 MB, 1 GB, 1.5 GB, 2 GB, and 10 GB^[32,33].

The default Hadoop defines the system’s default settings, whereas the proposed strategy defines the processes we have created. Table 5 reveals the proposed method of Hadoop and the response timing of process the data in the default Hadoop. The

strategy^[33,34] implementation is far better than Hadoop’s settings. Table 5 shows the contrast between the implementation times of Hadoop and Hadoop’s system. When these data’s magnitude remains small, there is not any difference between the implementation time. However, the execution times of the proposed procedure are reduced.

5 Conclusion

Works done under this research have produced amazing results, including the ability to schedule schedules, data placement inequality matrix, clustering before timing scrutiny, and reducing the frequency of repeat, mapping, and domestic dependency to prevent and stop responses. This experience demonstrates that by defining the process to deal with different usage cases, one can reduce the overall cost of computing and benefit from the use of distributed systems for fast splitting. When archiving or processing large amounts of data, developing solutions to compress data and search for relevant files in a compressed form could be beneficial. The methodology employed here is based on the practical process of sorting, iterating, managing queries, and finally producing the output. Assume that the data have been clustered by some powerful clustering algorithms. In this case, significant time and effort could be saved by improving the performance of Hadoop’s heterogeneous clusters for MapReduce job processing. The major challenges encountered during this project were the incompatibility of systems, different nodes, and security features provided by Linux and Hadoop. Working on dedicated servers and clusters could help achieve smoother experiences rather than dealing with compatibility issues at critical times. The underlying architecture and usage could be improved by working on higher versions of Hadoop to incorporate the best offerings and avoid being bogged down by minor issues, such as compatibility.

References

- [1] M. S. Mahmud, J. Z. Huang, S. Salloum, T. Z. Emara, and K. Sadatdiyov, A survey of data partitioning and sampling methods to support big data analysis, *Big Data Mining and Analytics*, vol. 3, no. 2, pp. 85–101, 2020.
- [2] M. D. Li, H. Z. Wang, and J. Z. Li, Mining conditional functional dependency rules on big data, *Big Data Mining and Analytics*, vol. 3, no. 1, pp. 68–84, 2020.
- [3] S. Salloum, J. Z. Huang, and Y. L. He, Random sample partition: A distributed data model for big data analysis, *IEEE Trans. Industr. Inform.*, vol. 15, no. 11, pp. 5846–5854, 2019.

- [4] R. H. Lin, Z. Z. Ye, H. Wang, and B. D. Wu, Chronic diseases and health monitoring big data: A survey, *IEEE Rev. Biomed. Eng.*, vol. 11, pp. 275–288, 2018.
- [5] Y. N. Tang, H. X. Guo, T. T. Yuan, Q. Wu, X. Li, C. Wang, X. Gao, and J. Wu, OEHadoop: Accelerate Hadoop applications by co-designing Hadoop with data center network, *IEEE Access*, vol. 6, pp. 25849–25860, 2018.
- [6] X. C. Hua, M. C. Huang, and P. Liu, Hadoop configuration tuning with ensemble modeling and metaheuristic optimization, *IEEE Access*, vol. 6, pp. 44161–44174, 2018.
- [7] D. Z. Cheng, X. B. Zhou, P. Lama, M. K. Ji, and C. J. Jiang, Energy efficiency aware task assignment with DVFS in heterogeneous Hadoop clusters, *IEEE Trans. Parallel Distrib. Syst.*, vol. 29, no. 1, pp. 70–82, 2018.
- [8] A. Kumar, A. Kumar, A. K. Bashir, M. Rashid, V. D. A. Kumar, and R. Kharel, Distance based pattern driven mining for outlier detection in high dimensional big dataset, *ACM Trans. Manag. Inf. Syst.*, vol. 13, no. 1, pp. 1–17, 2022.
- [9] A. Khaleel and H. Al-Raweshidy, Optimization of computing and networking resources of a Hadoop cluster based on software defined network, *IEEE Access*, vol. 6, pp. 61351–61365, 2018.
- [10] M. Malik, K. Neshatpour, S. Rafatirad, and H. Homayoun, Hadoop workloads characterization for performance and energy efficiency optimizations on microservers, *IEEE Trans. Multi-Scale Comput. Syst.*, vol. 4, no. 3, pp. 355–368, 2018.
- [11] Y. Yao, J. Y. Wang, B. Sheng, C. C. Tan, and N. F. Mi, Self-adjusting slot configurations for homogeneous and heterogeneous Hadoop clusters, *IEEE Trans. Cloud Comput.*, vol. 5, no. 2, pp. 344–357, 2017.
- [12] H. Alshammari, J. Lee, and H. Bajwa, H2Hadoop: Improving Hadoop performance using the metadata of related jobs, *IEEE Trans. Cloud Comput.*, vol. 6, no. 4, pp. 1031–1040, 2018.
- [13] I. Ullah, M. S. Khan, M. Amir, J. Kim, and S. M. Kim, LSTPD: Least slack time-based preemptive deadline constraint scheduler for Hadoop clusters, *IEEE Access*, vol. 8, pp. 111751–111762, 2020.
- [14] R. R. Parmar, S. Roy, D. Bhattacharyya, S. K. Bandyopadhyay, and T. H. Kim, Large-scale encryption in the Hadoop environment: Challenges and solutions, *IEEE Access*, vol. 5, pp. 7156–7163, 2017.
- [15] S. Kumar and M. Singh, A novel clustering technique for efficient clustering of big data in Hadoop ecosystem, *Big Data Mining and Analytics*, vol. 2, no. 4, pp. 240–247, 2019.
- [16] W. Huang, L. K. Meng, D. Y. Zhang, and W. Zhang, In-memory parallel processing of massive remotely sensed data using an Apache spark on Hadoop YARN model, *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, vol. 10, no. 1, pp. 3–19, 2017.
- [17] M. Soualhia, F. Khomh, and S. Tahar, A dynamic and failure-aware task scheduling framework for Hadoop, *IEEE Trans. Cloud Comput.*, vol. 8, no. 2, pp. 553–569, 2020.
- [18] D. Tao, Z. W. Lin, and B. X. Wang, Load feedback-based resource scheduling and dynamic migration-based data locality for virtual Hadoop clusters in OpenStack-based clouds, *Tsinghua Science and Technology*, vol. 22, no. 2, pp. 149–159, 2017.
- [19] P. Qin, B. Dai, B. X. Wang, and G. Xu, Bandwidth-aware scheduling with SDN in Hadoop: A new trend for big data, *IEEE Syst. J.*, vol. 11, no. 4, pp. 2337–2344, 2017.
- [20] X. Y. Wang, M. Veeraraghavan, and H. Y. Shen, Evaluation study of a proposed Hadoop for data center networks incorporating optical circuit switches, *J. Opt. Commun. Netw.*, vol. 10, no. 8, pp. C50–C63, 2018.
- [21] Y. Q. Chen, Y. Zhou, S. Taneja, X. Qin, and J. Z. Huang, aHDFS: An erasure-coded data archival system for Hadoop clusters, *IEEE Trans. Parallel Distrib. Syst.*, vol. 28, no. 11, pp. 3060–3073, 2017.
- [22] Z. Z. Li, H. Y. Shen, W. Ligon, and J. Denton, An exploration of designing a hybrid scale-up/out Hadoop architecture based on performance measurements, *IEEE Trans. Parallel Distrib. Syst.*, vol. 28, no. 2, pp. 386–400, 2017.
- [23] H. F. Wang and Y. P. Cao, An energy efficiency optimization and control model for Hadoop clusters, *IEEE Access*, vol. 7, pp. 40534–40549, 2019.
- [24] N. M. F. Qureshi, D. R. Shin, I. F. Siddiqui, and B. S. Chowdhry, Storage-tag-aware scheduler for Hadoop cluster, *IEEE Access*, vol. 5, pp. 13742–13755, 2017.
- [25] Z. Z. Li and H. Y. Shen, Measuring scale-up and scale-out Hadoop with remote and local file systems and selecting the best platform, *IEEE Trans. Parallel Distrib. Syst.*, vol. 28, no. 11, pp. 3201–3214, 2017.
- [26] Y. P. Zheng and G. Y. Chen, Energy analysis and application of data mining algorithms for internet of things based on Hadoop cloud platform, *IEEE Access*, vol. 7, pp. 183195–183206, 2019.
- [27] C. T. Chen, L. J. Hung, S. Y. Hsieh, R. Buyya, and A. Y. Zomaya, Heterogeneous job allocation scheduler for Hadoop mapreduce using dynamic grouping integrated neighboring search, *IEEE Trans. Cloud Comput.*, vol. 8, no. 1, pp. 193–206, 2020.
- [28] P. Q. Jin, X. J. Hao, X. L. Wang, and L. H. Yue, Energy-efficient task scheduling for CPU-intensive streaming jobs on Hadoop, *IEEE Trans. Parallel Distrib. Syst.*, vol. 30, no. 6, pp. 1298–1311, 2019.
- [29] K. Sridharan, G. Komarasamy, and S. Daniel Madan Raja, Hadoop framework for efficient sentiment classification using trees, *IET Netw.*, vol. 9, no. 5, pp. 223–228, 2020.
- [30] Z. C. Dou, I. Khalil, A. Khreishah, and A. Al-Fuqaha, Robust insider attacks countermeasure for Hadoop: Design and implementation, *IEEE Syst. J.*, vol. 12, no. 2, pp. 1874–1885, 2018.
- [31] R. Agarwal, A. S. Jalal, and K. V. Arya, Local binary hexagonal extrema pattern (LBH_XEP): A new feature descriptor for fake iris detection, *Vis. Comput.*, vol. 37, no. 6, pp. 1357–1368, 2021.

- [32] R. Agarwal, A. S. Jalal, and K. V. Arya, Enhanced binary hexagonal extrema pattern (EBH_XEP) descriptor for iris liveness detection, *Wirel. Pers. Commun.*, vol. 115, no. 3, pp. 2627–2643, 2020.
- [33] R. Agarwal, A. S. Jalal, and K. V. Arya, A multimodal liveness detection using statistical texture features and spatial analysis, *Multimed. Tools Appl.*, vol. 79, no. 19, pp. 13621–13645, 2020.
- [34] R. Agrawal, A. S. Jalal, and K. V. Arya, Fake fingerprint liveness detection based on micro and macro features, *Int. J. Biom.*, vol. 11, no. 2, pp. 177–206, 2019.



Ankit Kumar received the BEng degree in computer science & engineering from West Bengal Technical University, Kolkata, India in 2010, the MEng degree in computer science engineering from Indian Institute of Information Technology, Allahabad (IIT Allahabad), India in 2012, and the PhD degree in computer science from Sri Satya

Sai University of Technology & Medical Sciences, Sehore, India in 2022. Currently he is working as an assistant professor at Department of Computer Engineering and Application, GLA University, Mathura. He has developed optimization algorithms in machine learning and data science. His current research interests include machine learning, deep learning, big data, evolutionary computation and its application in real-world, and optimization problems especially in optimization medical applications. He has published more than 65 research papers in reputed journals and conferences in high indexing databases and has 9 patents granted from Australia and India. He has authored 2 edited books published by Apple Publisher and CRC press. He has completed 1 funded research project from the TEQIP-3. He is an associate editor for reputed journals and publishers.



Surbhi Bhatiya has ten years of rich teaching and academic experience. She has earned professional management certification from PMI, USA. She is currently working as an assistant professor at the Department of Information Systems, College of Computer Sciences and Information Technology, King Faisal

University, Saudi Arabia. She is also an adjunct professor at Shoolini University, Himachal Pradesh, India. She is an associate editor for reputed journals and publishers. She has published more than 70 research papers in reputed journals and conferences in high indexing databases, and has 9 patents granted from USA, Australia, and India. She has authored 3 solo books and 9 edited books published by Springer, Wiley, and Elsevier. She has completed 5 funded research projects from the Deanship of Scientific Research, King Faisal University, and the Ministry of Education, Saudi Arabia. Her research interests include information systems, and data analytics and sentiment analysis.



Neeraj Varshney is presently working as an assistant professor at the Department of Computer Engineering and Application, GLA University, Mathura. He has done MEng and MCA degrees and is pursuing the PhD degree. He has vast experience in the teaching and research of computer science. He has published more than 20

research papers in reputed international and Indian journals. His research areas include vehicular networks, digital image processing, machine learning, and time series data analysis.



Kamred Udham Singh received the PhD degree from Banaras Hindu University, India in 2019. From 2015 to 2016, he was a junior research fellow, and from 2017 to 2019, he was a senior research fellow with UGC (University Grant Commission), India. In 2019, he became an assistant professor at the School of Computing, Graphic Era Hill

University, India. He is with the School of Computing, Graphic Era Hill University, India. His research interests include image security and authentication, deep learning, medical image watermarking, and information security. He has published several research papers in international peer-reviewed journals. He contributes his expertise as a reviewer and editor in many reputed journals.