

XHAMI – extended HDFS and MapReduce interface for Big Data image processing applications in cloud computing environments

Raghavendra Kune^{1,*}, Pramod Kumar Konugurthi¹, Arun Agarwal²,
Raghavendra Rao Chillarige² and Rajkumar Buyya³

¹*Advanced Data Processing Research Institute, Department of Space, Hyderabad, India*

²*School of Computer and Information Sciences, University of Hyderabad, Hyderabad, India*

³*Cloud Computing and Distributed Systems (CLOUDS) Lab, Department of Computing and Information Systems, University of Melbourne, Melbourne, Vic. Australia*

SUMMARY

Hadoop distributed file system (HDFS) and MapReduce model have become popular technologies for large-scale data organization and analysis. Existing model of data organization and processing in Hadoop using HDFS and MapReduce are ideally tailored for search and data parallel applications, for which there is no need of data dependency with its neighboring/adjacent data. However, many scientific applications such as image mining, data mining, knowledge data mining, and satellite image processing are dependent on adjacent data for processing and analysis. In this paper, we identify the requirements of the overlapped data organization and propose a two-phase extension to HDFS and MapReduce programming model, called XHAMI, to address them. The extended interfaces are presented as APIs and implemented in the context of image processing application domain. We demonstrated effectiveness of XHAMI through case studies of image processing functions along with the results. Although XHAMI has little overhead in data storage and input/output operations, it greatly enhances the system performance and simplifies the application development process. Our proposed system, XHAMI, works without any changes for the existing MapReduce models and can be utilized by many applications where there is a requirement of overlapped data. Copyright © 2016 John Wiley & Sons, Ltd.

Received 3 November 2015; Revised 25 March 2016; Accepted 6 June 2016

KEY WORDS: cloud computing; Big Data; Hadoop; MapReduce; extended MapReduce; XHAMI; image processing; scientific computing; remote sensing

1. INTRODUCTION

The amount of textual and multimedia data has grown considerably large in recent years because of the growth of social networking, healthcare applications, surveillance systems, earth observation sensors, and so on. This huge volume of data in the world has created a new field in data processing called as Big Data 1, which refers to an emerging data science paradigm of multidimensional information mining for scientific discovery and business analytics over large-scale scalable infrastructure 2. Big Data handles massive amounts of data collected over time, which is otherwise difficult task to analyze and handle using common database management tools 3. Big Data can yield extremely useful information, however apart, urges new challenges both in data organization and in processing 4.

Hadoop 5 is an open-source framework for storing, processing, and analysis of large amounts of distributed semi-structured/unstructured data 6. The origin of this framework comes from Internet

*Correspondence to: Raghavendra Kune, Department of Space, Advanced Data Processing Research Institute, Hyderabad 500009, India.

†E-mail: raghav.es@gmail.com

search companies like Yahoo and Google, who needed new processing tools and models for web page indexing and searching. This framework is designed for data parallel processing at Petabyte and Exabyte scales distributed on the commodity computing nodes. Hadoop cluster is a highly scalable architecture that spawns both computing and data storage nodes horizontally for preserving and processing large-scale data to achieve high reliability and high throughput. Therefore, Hadoop framework and its core subcomponents – HDFS 7, 8 and MapReduce (MR) 9–11 – are gaining popularity in addressing several large-scale applications of data-intensive computing in several domain-specific areas like social networking, business intelligence, and scientific analytics for analyzing large-scale, rapidly growing, variety structures of data.

The advantages of HDFS and MR in Hadoop ecosystem are horizontal scalability, low-cost setup with commodity hardware, ability to process semi-structured/unstructured data, and simplicity in programming. However, HDFS and MR offer tremendous potential for gaining maximum performance but, because of its certain inherent limiting features, are unable to support applications with overlapped data processing requirements. In the following sections, we describe one such domain-specific application in remote sensing (RS) image processing and their requirements.

1.1. Remote sensing imaging applications

Earth observation satellite sensors provide high-resolution satellite imagery having image scene sizes from several megabytes to gigabytes. High-resolution satellite imagery, for example, Quick Bird, IKONOS, Worldview, and Indian Remote Sensing CARTOSAT 12, is used in various applications of information extraction and analysis in domains such as oil/gas mining, engineering construction like 3D urban/terrain mapping, GIS developments, defense and security, environmental monitoring, media and entertainment, and agricultural and natural resource exploration. Because of increase in the numbers of satellites and technology advancements in the RS, both the data sizes and their volumes are increasing on a daily basis. Hence, organization and analysis of such data for intrinsic information is a major challenge.

Ma *et al.* 13 discussed challenges and opportunities in RS Big Data computing and focussed on RS data-intensive problems, analysis of RS Big Data, and several techniques for processing RS Big Data. Two-dimensional structured representation of images, and majority of the functions in image processing being highly parallelizable, the HDFS way of organizing the data as blocks and usage of MR functions for processing each block as independent map function, makes Hadoop a suitable platform for large-scale high-volume image processing applications.

1.2. Image representation

An image is a two-dimensional function $f(x,y)$ as depicted in Figure 1, where x and y are spatial (plane) coordinates, and the amplitude of ' f ' at any pair of coordinates (x,y) is called intensity or gray level of the image at that point 14. Image data mining is a technology that aims in finding useful information and knowledge from large-scale image data 15 until a pattern in the image becomes obvious. This involves the usage of several image processing techniques such as enhancement, classification, segmentation, and object detection, which could use in turn several combinations of linear/morphological spatial filters 14 to achieve the result. Image mining is the process of searching and discovering valuable information and knowledge in large volumes of data. Image mining follows basic principles from concepts in databases, machine learning, statistics, pattern recognition, and soft computing. Image mining is an emerging research field in geosciences due to the large-scale increase of data that lead to new promising applications. For example, the use of very high-resolution satellite images in earth observation systems now enables the observation of small objects, while the use of very high temporal resolution images enables monitoring of changes at high frequency for detecting the objects.

However, actual data analysis in geosciences or earth observation techniques suffers from the huge amount of complex data to process. Indeed, earth observation data (acquired from optical, radar, and hyper spectral sensors installed on terrestrial, airborne, or space-borne platforms) is often heterogeneous, multi-scale, and composed of complex objects. Segmentation algorithms,

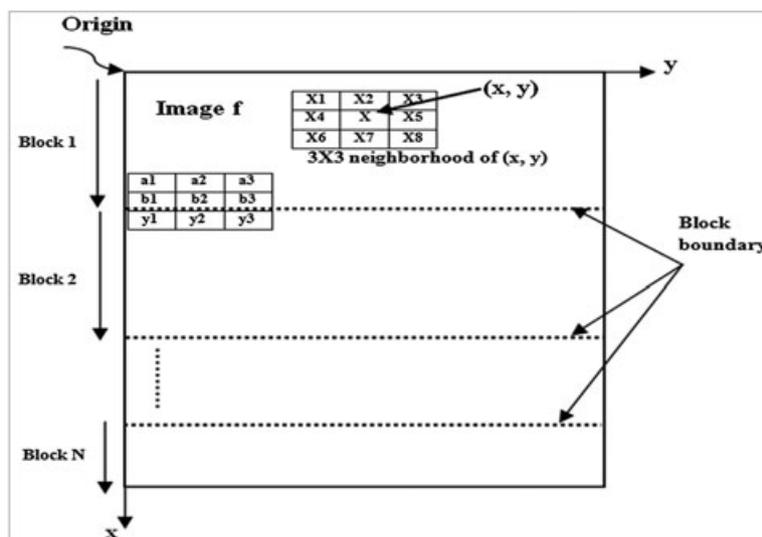


Figure 1. Image representation with segmented blocks.

unsupervised and supervised classification methods, descriptive and predictive spatial models, and algorithms for large time series would be applied to assist experts in their knowledge discovery.

Many of the algorithms applied during image mining techniques such as linear/morphological spatial filters demand use of adjacent pixels for processing the current pixel. For example, as depicted in Figure 1, a smoothing operation performs weighted average of a 3×3 kernel window; hence, the output of pixel X depends on the values of $X1$, $X2$, $X3$, $X4$, $X5$, $X6$, $X7$, and $X8$. If an image is organized as several segmented blocks, then the pixels on the edge boundaries of the segmented blocks would require the pixels from their next immediate adjacent block, while processing. The Image with segmented blocks and block boundaries representation is depicted in Figure 1. Therefore, if the image is organized as several physical distributed data blocks, image filtering operations cannot be performed on the edge pixels, unless overlap data required are preserved in the blocks. For example, processing the pixels such as $b1$, $b2$, and $b3$ of block 1 requires its adjacent block 2 pixels such as $y1$, $y2$, and $y3$. It is easy to obtain the overlapping pixels, for an image organized as a single large file, by moving the file pointer back and forth to read the data of interest. But if the images are organized as segmented blocks, then the application demands the overlapping data at data nodes for processing, which in turn requires moving of this overlapped data to the data nodes, which accounts for large I/O overheads as the data blocks are distributed at several nodes.

Hadoop and several other similar application implementations (Section 2) split the data based on a fixed size, which results in partitioning of data as shown in Figure 1. Each of the blocks is written to different data nodes. Therefore, the boundary pixels of entire line $b1$, $b2$, $b3$... in each block cannot be processed, as the adjacent pixels are not available at the respective data nodes. Similarly, the pixels marked as $y1$, $y2$, $y3$, $y4$,... cannot be performed straight away. To process these boundary pixels, that is, the start line and end line in each block, a customized map function to read additional pixels from a different data node is essential; otherwise, the output would be incorrect. These additional read operations for each block increase I/O overhead significantly.

1.3. Our contributions

To meet the requirements of applications with overlapped data, we propose an extended HDFS and MR interface, called XHAMI, which offers a two-phase extension to HDFS and MR programming model. The extended interfaces are presented as APIs and implemented in the context of image processing application domain. We demonstrated effectiveness of XHAMI through case studies of image processing functions along with the results. Our experimental results reveal that XHAMI greatly enhances the system performance and simplifies the application development process. It works without any changes for the existing MR models, and hence, it can be utilized by many applications where there is a requirement of overlapped data.

1.4. Paper organization

The rest of the paper is organized as follows. Section 2 describes related work in image processing with HDFS and MR over Hadoop framework. Section 3 describes XHAMI system for large-scale image processing applications as a two-phase extension over the conventional Hadoop HDFS and MR system. The first extension discusses the data organization over HDFS, and the second illustrates high-level packages for image processing. XHAMI hides the low-level details of the data organization over HDFS and offers a high-level APIs for processing large-scale images organized in HDFS using MR computations. Section 4 describes experimental results of XHAMI compared with conventional Hadoop for data organization, followed by customized Hadoop MR, and other related systems like Hadoop Image Processing Interface (HIPI) 16 for MR computations. Section 5 presents conclusions and future work.

2. RELATED WORK

Image processing and computer vision algorithms can be applied as multiple independent tasks on large-scale data sets simultaneously in parallel on a distributed system to achieve higher throughputs. Hadoop 5 is an open-source framework for addressing large-scale data analytics uses HDFS for data organization and MR as programming models. In addition to Hadoop, there are several other frameworks like Twister 17 for iterative computing of streaming text analytics, and Phoenix 18 used for map and reduce functions for distributed data-intensive message passing interface kind of applications.

Kennedy *et al.* 19 demonstrated the use of MR for labeling 19.6 million images using nearest neighbor method. Shi *et al.* 20 presented use of MR for content-based image retrieval and discussed the results obtained by using around 400,000 images approximately. Yang *et al.* 21 presented a system MIFAS for fast and efficient access to medical images using Hadoop and Cloud computing. Kocalkulak *et al.* 22 proposed a Hadoop-based system for pattern image processing of intercontinental missiles for finding the bullet patterns. Almeer *et al.* 23 designed and implemented a system for RS image processing with the help of Hadoop and Cloud computing systems for small-scale images. Demir 24 *et al.* discussed the usage of Hadoop for small size face detection images. All these systems describe the bulk processing of small size images in batch mode over HDFS, where each map function processes the complete image.

White *et al.* 25 discussed the overheads that can be caused because of small size files, which are considerably smaller than the block size in HDFS. A similar approach is presented by Sweeney *et al.* 16 and presented HIPI as an extension of MR APIs for image processing applications. HIPI operates on the smaller image files and bundles the data files (images) into a large single data file called HIPI Image Bundle (HIB), and the indexes of these files are organized in the index file.

Potisepp 26 discussed the processing small/regular images of total 48,675 by aggregating them into large data set and processed them on Hadoop using MR as sequential files, similar to the one addressed by HIPI, and also presented feasibility study as a proof-of-concept test for a single large image as blocks and overlapping pixels for non-iterative algorithms image processing. However, no design, or solution, or methodology has been suggested to either to Hadoop or to MR for either image processing applications or for any other domain, so that the methodology works for existing as well as new models that are under consideration.

Papers discussed earlier demonstrated the usage of Hadoop for image processing applications and compared the performance of image processing operations over single PC system running on a conventional file system versus Hadoop cluster. Few other papers have demonstrated the extension of Hadoop, called HIPI to solve the small files problem, which combine the smaller images into large bundle and process them using HIPI Bundle (HIB). These systems had limitations in addressing the spatial image filters applications as the overlap data are not present among the adjacent blocks for processing. In this paper, we address the issues related to organization and processing of single large volume images, which are in general in data sizes ranging from megabytes to gigabytes, and their processing using MR. To address the issues, we discuss the extensions proposed of the Hadoop for HDFS and MR and present those extensions XHAMI. Although XHAMI demonstrates

RS/geosciences data processing, the same can be used for other domains where data dependencies are major requirements, for example, biomedical imaging.

3. XHAMI – EXTENDED HDFS AND MAPREDUCE

In this section, we describe XHAMI – the extended software package of Hadoop for large-scale image processing/mining applications. First, we present XHAMI APIs for reading and writing (I/O), followed by MR for distributed processing. We discuss two sample case studies, that is, histogram and image smoothening operations. Histogram computes the frequency of pixel intensity values in the image, and smoothening operation uses spatial filters like Sobel and Laplacian 14. Later, we discuss how XHAMI can be used for data organization, designing user-specific image-related MR functions and extending the functionality for other image domain-specific applications.

3.1. XHAMI – HDFS I/O extensions for domain-specific applications

Figure 2 depicts the sequence of steps in reading/writing the images using XHAMI software library over Hadoop framework. Initially, client uses XHAMI I/O functions (step 1) for reading or writing the data. The client request is translated into create() or open() by XHAMI and sent to distributed file system (step 2). Distributed file system instance calls the name node to determine the data block locations (step 3). For each block, the name node returns the addresses of the data nodes for writing or reading the data. Distributed file system returns FSDataInput/Output Stream, which in turn will be used by XHAMI to read/write the data to/from the data nodes. XHAMI checks file format, if the format is in image type (step 4), then metadata information such as file name, total scans, total pixels, total numbers of bands in the image, and the number of bytes per pixel is stored in HBASE 28, and this simplifies header information reading as and when required through HBASE queries; otherwise, reading the header block by block is tedious and time-consuming process.

Later on, XHAMI calls FSDataInput/Output Stream either to read/write the data to/from the respective data nodes (step 5). Steps 6 and 7 are based on standard HDFS data reading/writing in the pipelining way. Each block is written with the header information corresponding to the blocks, that is, blockid, start scan, end scan, overlap scan lines in the block, scan length, and size of the

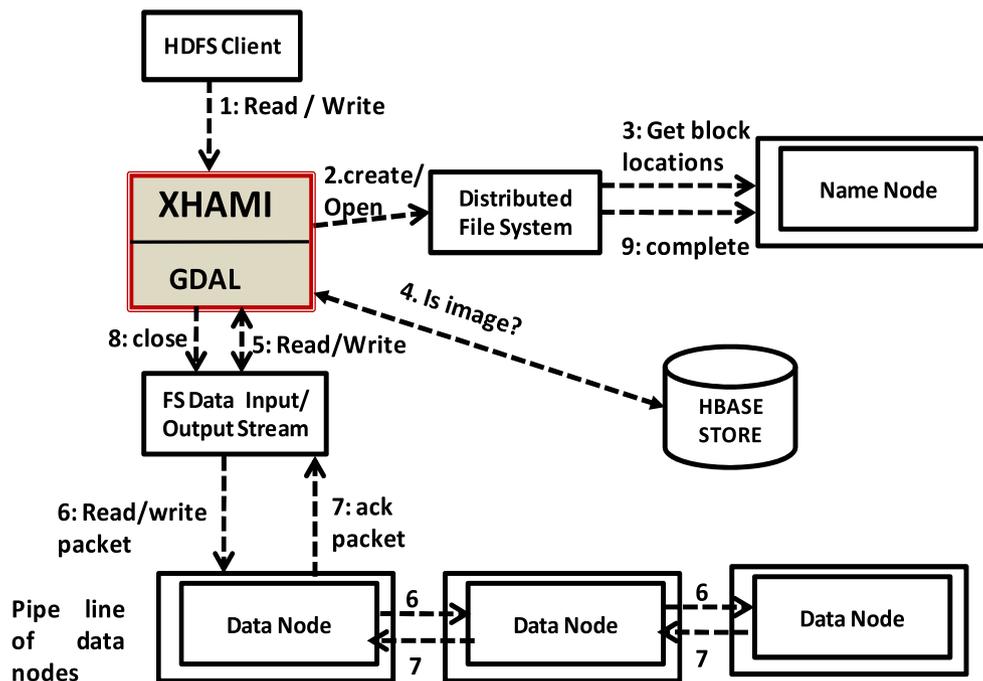


Figure 2. XHAMI for read/write operations.

block. Finally, after the read/write operation, the request is made for closing the file (step 8), and the status (step 9) is forwarded to the name node. In the succeeding sections, we describe techniques developed for data organization followed by extended APIs for HDFS and MR.

3.2. Data organization

Single large image data are organized as blocks, with an overlap with its next immediate blocks. The segmented techniques used for data organization are shown in Figure 3. Figure 3a depicts an image as a one-dimensional sequence of bytes, Figure 3b depicts block segmentation in horizontal direction, and Figure 3c shows the organization in both horizontal and vertical directions.

Image blocks can be constructed in two ways, that is, (i) unidirectional: partitioning across the scan line direction as shown in Figure 3b and (ii) bidirectional: partitioning both horizontal and vertical directions as shown in Figure 3c. While construction, it is essential to ensure that no split takes place within the pixel byte boundaries. The methods are described as follows.

- (i) Unidirectional split: blocks are constructed by segmenting the data in across scan line (horizontal) direction. Each block is written with the additional lines at the end of the block.
- (ii) Bidirectional split: splitting the file into blocks in both horizontal and vertical directions. The split results in the blocks, for which the first and last blocks have overlapped with their adjacent two blocks, and all the remaining blocks have overlapped with their adjacent four blocks. This type of segmentation results in large storage overhead that is approximately double the size of the unidirectional segment construction. This type of organization is preferred while images have larger scan line lengths.

In the current version of XHAMI package, data organization is addressed for unidirectional segmented blocks; however, it can be extended for bidirectional split. The segmentation procedure is described as follows.

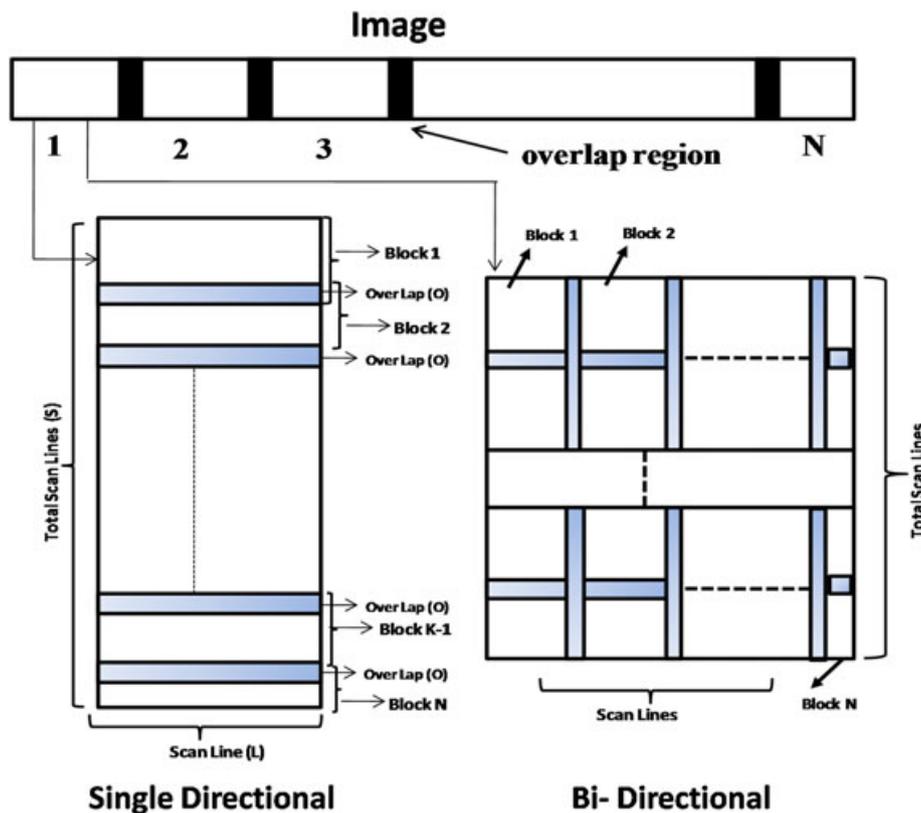


Figure 3. Block construction methods.

Scan lines for each block Sb are computed as

$$Sb = \lceil H / (L * P) \rceil$$

H = HDFS default block length in Mbytes.

L = length of scan line, that is, total pixels in the scan line.

P = pixel length in bytes.

S = total number of scan lines.

The total number of blocks T , having overlap of α number of scan lines, is

$$T = S / Sb$$

If $T * \alpha > Sb$, then $T = T + 1$.

The start and end scan lines $B_{i,s}$ and $B_{i,e}$ in each block are given as follows: N representing total scans in the image.

$$B_{i,s} = \begin{cases} 1, & i = 1 \\ B_{i-1,e-\alpha+1}, & 1 < i < T \\ B_{N-1,e-\alpha+1} & i = T \end{cases}$$

$$B_{i,e} = \begin{cases} B_{i,s} + Sb - 1 & 1 \leq i < T \\ Sb i & i = T \end{cases}$$

Block length is computed as follows.

$$R_i = (B_{i,e} - B_{i,s} + 1) * L * P, \quad 1 \leq i \leq T$$

The blocks are constructed with metadata information in the header, such as blockid, start scan, end scan, overlap scan lines in the block, scan length, and block length. Though, metadata adds small additional storage overheads but simplifies the processing activity during Map phase, for obtaining the total number of pixels, number of bands, bytes per pixel, and so on, and also helps to organize the blocks in the order during the combine/merge phase using blockid.

3.3. XHAMI package description

XHAMI offers Software Development Kit for Hadoop-based large-scale domain-specific data-intensive applications designing. It provides high-level packages for data organization and for MR-based processing simplifying the development and quick application designing by hiding several low-level details of image organization and processing. XHAMI package description is as follows: XhamiFileIOFormat is the base class for all domain-specific applications, which is placed under the package xhami.io, as shown in Figure 4.

XhamiFileIOFormat extends FileInputFormat from the standard Hadoop package, and the implementation of several methods, for hiding the low-level handling of data for HDFS data organization, is handled by this package. The major methods offered under XhamiFileIOFormat class are (i) setting the overlap; (ii) getting the overlap; and (iii) reading, writing, and seeking the data without knowing/knowledge of the low-level details of the data.

XhamiFileIOFormat is used as base class for several application developers for file I/O functionality and implements further for domain-specific operations. Xhami.io.XhamiImage is an abstract class that provides the methods for the implementation of image processing domain-specific functionality by extending XhamiFileIOFormat. XhamiImage extends XhamiFileIOFormat class and implements several image processing methods for setting the header/metadata information of the

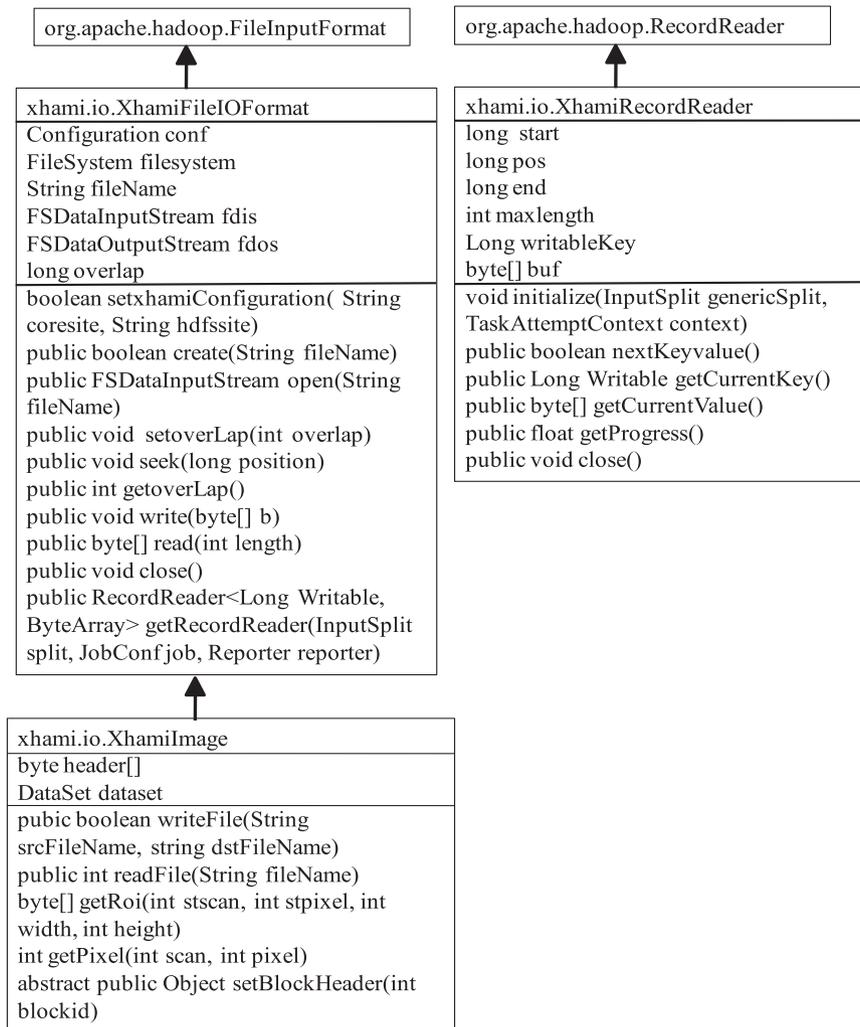


Figure 4. XHAMI I/O package.

image, blockwise metadata information, and for reading/writing the image blocks into the similar format of that original file, using Geographical Data Abstraction Layer (GDAL) library 27, and offers several methods such as reading the region of interest, getting the scan lines, pixel, and reading the pixel gray value, by hiding the low-level details of file I/O. Several methods in XhamiFileIOFormat, XhamiImage classes, and XhamiRecordReader classes are described in Tables I–III, respectively. XhamiImage implements GDAL functionality using the DataSet object for image-related operations. XhamiImage class can be extended by the Hadoop-based image processing application developer by setting up their own implementation of setBlockHeader method. XhamiRecordReader reads the buffer data and sends to the Map function for processing.

Package hierarchy of MR for image processing domain is shown in Figure 5, and methods in the packages are described in Table IV.

In what follows, we describe sample implementation of XHAMI MR functions for image processing domain.

3.4. XHAMI – MapReduce functions

In this section, we describe the extensions for Map and Reduce functions for image processing applications. Based on the image processing operation, either map function alone or both map and reduce functions are implemented. For example, edge detection operation does not require the

Table I. Description of methods in XhamiFileIOFormat class.

Method name	Method description	Return value
boolean setxhamiConfiguration(String coressite, String hdfssite)	Sets the HDFS configuration parameters such as coressite and hdfssite. This function in turn uses the Configuration object and calls addResource methods of its base class, to set establish the connectivity to HDFS site.	If the configuration parameters are correct, then boolean value true is returned. In case wrong supply of arguments, or if the parameter files are not available, or because of invalid credentials, or else HDFS site may be down, false will be returned.
boolean create(String fileName)	Create the file with name filename to write to HDFS. Checks if the file already exists. This function is used before the file is to be written to HDFS.	Returns true if the file is not present in HDFS, or else returns false.
FSDatInputStream open(String filename)	Checks if the file is present in the HDFS.	If the file is present, FSDatInputStream having the object value is returned; otherwise, FSDatInputStream with value having null is returned.
void setoverLap(int overlap)	Used to set the overlap across the segmented blocks. The supplied overlap value is an integer value corresponding to the overlap size in bytes.	—
void seek(long position)	Moves the file pointer to the location position in the file.	—
int getoverLap()	Reads the overlap value set for the file while writing to HDFS.	Return the overlap value if set for the file or else returns -1.
void write(byte[] b)	Writes b number of bytes to the file.	—
byte[] read(int length)	Reads length number of bytes from the file.	—
void close()	Closes the data pointers those are opened for reading/writing the data.	Returns the data in the byte array format.
RecordReader<Long Writable, ByteArray> getRecordReader(InputSplit split, JobConf job, Reporter reporter)	Reads the Xhami compatible record reader in bytes, for MapReduce computing by overriding the RecordReader method of FileInputFormat class. The compatible here means the window size to be read for processing the binary image for processing. This would be supplied as argument value to the Image Processing MapReduce function. If the value is of type fixed, then the entire Block is read during processing.	Returns the Default Hadoop RecordReader Object for processing by the MapReduce job.

Table II. XhamiImageclass description.

Method name	Method description	Return value
boolean writeFile(String srcFileName, String dstFileName)	Used for writing the contents of the file srcFileName, to the destination dstFileName.	Boolean value true is returned if the writing is successful, or else false is returned.
int readfile(String fileName)	Set the file fileName to read from the HDFS. Returns the total numbers of blocks that the file is organized in HDFS.	Number of blocks that the file is stored in HDFS. If the file does not exist, -1 is returned.
byte[] getRoi(int stscan, int stpixel, int width, int height)	Reads the array of bytes from the file already set, starting at stscan and pixel, with a block of size width and height bytes.	If successful, returns byte array read, or else returns NULL object.
int getPixel(int scan, int pixel)	Reads the pixel value at the location scan and pixel.	Returns pixel(gray) value as integer.
abstract public Object setBlockHeader(int blockid)	Abstract method that would be overwritten by XHAMI application developers for image processing domain applications.	Header information type casted to Object data type.

Table III. Methods in XhamiRecordReader.

Method name	Method description	Return value
Initialize(InputSplit genericSplit, TaskAttemptContext context)	Overrides the Initialize method of standard Hadoop Record reader method. Implements the own split method, which reads the content that is compatible with the Xhami File data format, hiding the overlap block size.	—
boolean nextKeyvalue()	Overrides the nextKeyvalue method of its base class RecorReader.	Returns true if it can read the next record for the file, or else return false.
Long Writable getCurrentKey()	getCurrentKey method of its base class is overridden.	Return Writable Object of the record recorder method.
float getProgress()	Overriding method, to send the progress of the data read.	Return float false representing the percentage of data read so far from the XhamiRecordRecorder, corresponding to the InputSplit.
byte[] getCurrentValue()	Reads the bytes array to be sent for computing for Map function.	Return byte array if true, else returns NULL object.
void close()	Closes the record reader object.	—

reducer, as the resultant output of the map function is directly written to the disk. Each map function reads the block numbers and metadata of the corresponding blocks.

Read operations can be implemented in two ways in HDFS, one way is to implement own split function, ensuring the split does not happen across the boundaries, and the other one is to use FIXED LENGTH RECORD of FixedLengthInputFormat class. The package offers the implementations for both FIXED LENGTH RECORD and custom formatter XhamiRecordRecorder.

(1) MapReduce Sobel edge detection sample implementation

Edges that characterize boundaries in images are areas with strong intensity contrasts – a jump in intensity from one pixel to the next. There are many ways to perform edge detection. However, the majority of different methods may be grouped into two categories: gradient and Laplacian. The gradient

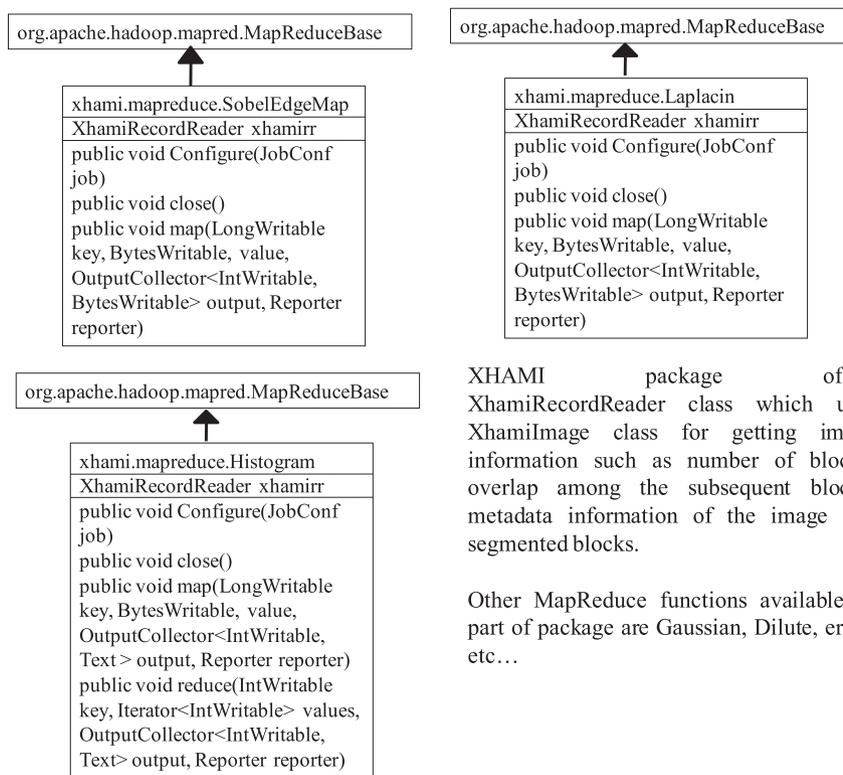


Figure 5. XHAMI MapReduce package for image processing domain.

Table IV. Description of XHAMI MapReduce classes for image processing domain.

MapReduce class	Description	Return value
Sobel	Implementation of Sobel spatial edge detection filter. It has map function implementation only, and the Reduce is not required, as the output of the map itself is directly written, as it does not require any collection of the map inputs for processing further. This implementation hides the several details such as overlapping pixels across the blocks and the kernel window size to be read for processing. Output is written to the HDFS file system.	Output images with the detected edges.
Laplacian	Implementation of Laplacian differential edge detection filter. It has map function implementation but not reduce method. Reduce is not required, as the output of the map itself is directly written, as it does not require any collection of the map inputs for processing further. This implementation hides the several details such as overlapping pixels across the blocks and the kernel window size to be read for processing. Output is written to the HDFS file system.	Output images with the detected edges.
Histogram	Implements both Map and Reduce functions. Map collects the count of the pixel (gray) value, and reduce does the aggregation of the collected numbers from the map functions. While processing, it does not consider the overlapping pixels across the blocks.	Histogram of the image.

method detects the edges by looking for the maximum and minimum in the first derivative of the image. The Laplacian method searches for zero crossings in the second derivative of the image to find the edges. An edge has the one-dimensional shape of a ramp, and calculating the derivative of the image can highlight its location. In the map function, for edge detection, the combiner and reduce functions are not performed, as there is no need for aggregation of the individual map functions.

The differences between conventional Hadoop implementation and the XHAMI implementations are as follows: in the former, the data are organized as segmented as blocks, and there is no overlap of the line pixels across the blocks. Hence, it would be difficult to process the edge pixels of the blocks, and to process the edge pixels, one should obtain the two blocks and compute the overlap pixels before it is sent to the map function for processing. Also, it would be difficult to ensure that the split does not happen within the pixel while reading. But XHAMI hides all such low-level details of data organization such as lines or pixels overlap, no split within the pixels, number of blocks the image is organized as blocks, and header information of the blocks.

(2) MapReduce histogram sample implementation

Histogram operation computes frequency count of the pixel in the image. The histogram is computed as follows: first, the block and length of the block are read, and each block is mapped to one map function. The difference between the conventional implementation and the XHAMI MR histogram implementation is as follows: in the former, it is necessary to ensure that the split does not happen within the pixels. The later overcomes this problem by using XHAMI image implementation for data organization and ensures that overlap lines (pixels) are vomited during processing by the Map function.

3.5. Writing domain-specific applications by extending XHAMI package

XhamiFileIOFormat class is the base class, which hides the low-level details of data organization and processing for the several applications of binary data handling. This class offers several methods for reading, writing, seeking to the particular block of the image, and getting the overlap information among the subsequent blocks. XhamiImage class is an extended class of XhamiFileIOFormat, which offers several methods for handling the data for several applications in image processing domain. XhamiImage could be used for development of HDFS-based data organization readily, or else, one can extend the class XhamiFileIOFormat for handling similar kind of image processing domain applications of their own interest. In what follows, we describe the procedure for writing and reading the images in HDFS format using by extending XhamiImage for writing XHAMI-based applications.

- Extend XhamiImage class and implement setBlockHeader method.
- Define header class and the necessary data types for implementation of the image processing application.
- Implement setBlockHeader method using the FSDataInputStream available in XhamImage class as member variable.
- Set the overlap required among the adjacent blocks using setoverLap method.
- Assign the source file and destination files, using the writeFile method. Internally, this method computes the total file size, by computing the total numbers of blocks that the image gets divided into and writes the corresponding block header.
- The contents of the file using getBlockCount and getBlockData methods.

Table V shows a sample code describing how to extend XhamiImage class for writing the images along with the image-specific header while storing the image into HDFS.

4. PERFORMANCE EVALUATION

In this section, we present the experiments conducted for large size images of RS data having different dimensions (scans and pixels) and sizes varying approximately from 288 MB to 9.1 GB. First, we discuss data organization and I/O overheads for read and write, followed by performance comparison of image processing operations for histogram and Sobel edge detection filter. We conduct

Table V. XhamiImage extension for writing block header.

```

class ImageHeader implements Serializable{
    int blockid, startscan, endscan, overlapscan, scanlength, blocklength, bytesperpixel;
}
public class XhamiImageIOOperations extends XhamiImage{//other implementation specific to the application
@Override
public object setBlockHeader(int blockid){
ImageHeader ih= new ImageHeader();
//set the details and write to FSDataOutputStream
}
}

```

the experiments both on Hadoop using conventional APIs and on XHAMI libraries and discuss how XHAMI simplifies the programming complexity and increases the performance when applied to large-scale images. The experiments are conducted to analyze the two important factors: (i) performance gain/improvement of proposed XHAMI system versus Hadoop and similar systems, as discussed in Section 4.1, and (ii) I/O performance overheads for read and write operations, of XHAMI versus Hadoop-based HDFS data organization, discussed in Section 4.2.

For the experimental study, virtualization setup-based Xen hypervisor with a pool of four servers of Intel Xeon 64 bit architecture is used. The configuration of the nodes for RAM, disk storage capacity, and virtual CPUs for each node in the Hadoop cluster is shown in Table VI. Hadoop version 2.7 is configured in the fully distributed mode on all these virtual machines running with 64 bit ‘Cent OS’ operating system.

4.1. Performance comparisons of XHAMI, Hadoop (HDFS and MapReduce), and HIPI

Sample data sets used for experiments are from the Indian Remote Sensing satellite series, that is, CARTOSAT-1, and CARTOSAT-2A is shown in Table VII. In the table, the columns show that image size represents the original image size in bytes in regular file system, and the resulted image size indicates the size in bytes in HDFS with overlapping of five scan lines using the block length computation algorithm in unidirectional approach illustrated in Section 3.2. A sample image with overlap of five scan lines shown in red color is depicted in Figure 6. The results show a maximum of 0.25% increase in the image size, which is negligible.

Here, we discuss the I/O performance of HDFS data organization in XHAMI and default Hadoop for single large volume image handling, followed by MR image processing in XHAMI, customized

Table VI. System configuration.

Type	Processor type	Hostname	RAM (GB)	Disk (GB)
Name node	Intel Xeon 64 bit, 4 vCpus, 2.2 GHz	namenode	4	100
Job tracker	-do-	jobtracker	2	80
Data node 1	-do-	datanode1	2	140
Data node 2	Intel Xeon 64 bit, 4 vCpus, 2.2 GHz	datanode2	2	140
Data node 3	Intel Xeon 64 bit, 2 vCpus, 2.2 GHz	datanode3	2	140
Data node 4	-do-	datanode4	2	100

Table VII. Sample data sets used and the resultant image size.

Serial no.	Image size (in bytes)	Scan line length	Total scan lines	Resulted image size (in bytes)
1	288,000,000	12,000	12,000	288,480,000
2	470,400,000	12,000	19,600	471,240,000
3	839,976,000	12,000	34,999	841,416,000
4	1,324,661,556	17,103	38,726	1,327,911,126
5	3,355,344,000	12,000	139,806	3,361,224,000
6	9,194,543,112	6026	762,906	9,202,738,472

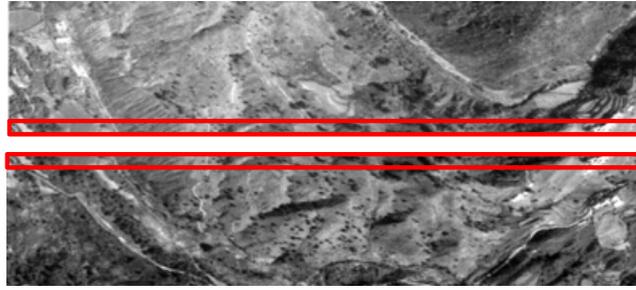


Figure 6. Image blocks with overlap highlighted in rectangular box.

MR for Hadoop, and HIPI. The data sets used for the experiments are shown in Table VII. Here, Hadoop, in its native form, cannot be used for MR image processing, because of the overlapped data requirements as discussed in Section 1; hence, we performed few customizations, to default Hadoop MR, like locating the neighbor adjacent block data and its location and migrating them to the node where Map functioning is to be computed. We present the results for image processing operations such as histogram and Sobel filter, followed by read and write overheads. The I/O overhead comparisons are presented for both Hadoop and XHAMI. The advantage of HIPI over customized Hadoop is the use of Java Image Processing Library. HIPI comes with processing the image formats like jpg and png files and hence does not require the additional implementation functions for image handling. As HIPI uses HIB formats to create a single bundle for smaller image files, but, the data sets used for the experiments are very large; hence, here HIB consists of a single image itself, unlike the smaller files bundled into a single image.

(a) Histogram operation

Histogram operation counts the frequency of the pixel intensity in the entire image, which is similar to counting the words in the file. The performance results of histogram operation for customized Hadoop MR, HIPI, and XHAMI are shown in Figure 7. For customized Hadoop MR, data are organized as non-overlapped blocks; for XHAMI, data blocks are with overlapping data, and extensions of HDFS and MR APIs are for processing. For HIPI processing, the data blocks are retrieved using HIB files, and for XHAMI, the data are retrieved from the data blocks having overlap region with their corresponding immediate subsequent block. The results show that all the three systems customized Hadoop, HIPI, and XHAMI performance are more or less similar and that XHAMI has little overhead, which is less than 0.8% with customized Hadoop, which is due to skipping of the overlapped scan lines while processing.

(b) Fixed mask convolution operation

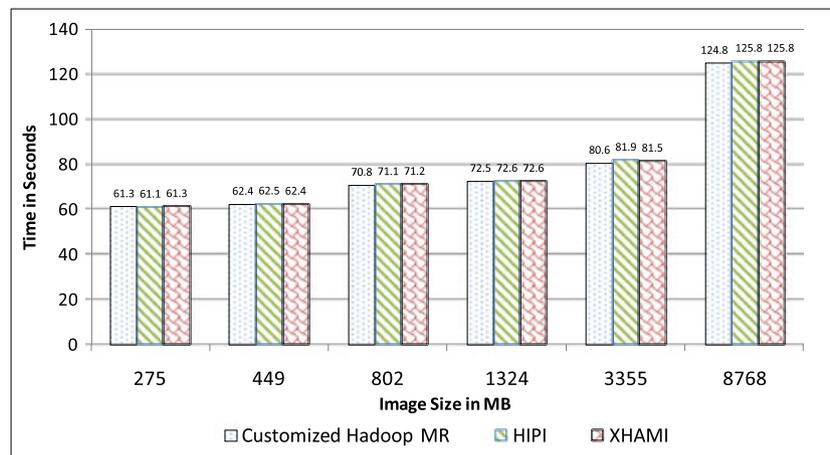


Figure 7. Histogram performance.

Convolution is a simple mathematical operation that is fundamental to many common image processing operators. The convolution is performed by sliding the kernel over the image, generally starting at the top left corner, so as to move the kernel through all the position where the kernel fits entirely within the boundaries of image. Convolution methods are the most common operations used in image processing, which uses the mask operator, that is, kernel for performing windowing operations on the images. Sobel operator is one of the commonly used methods for detecting edges in the image using convolution methods. In case if the image is organized as physical partitioned distributed blocks, then, convolution operations cannot process the edge pixels of such blocks, because of the non-availability of the adjacent blocks data on the same node. In conventional Hadoop-based HDFS and MR processing, the data are organized as physical partitioned blocks; hence, the edge pixels cannot be processed directly and demand the additional I/O overheads for processing the edge pixels of each block.

Here, we present the performance of the Sobel edge detection implementation in XHAMI and compare it with customized Hadoop MR and HIPI. In customized Hadoop MR, data are physically partitioned as non-overlapping data blocks, and for HIPI, data are organized as single large block for the file stored in HIB data format. For MR processing, customized Hadoop MR uses an additional functionality that is included in Map function for retrieving the adjacent block information corresponding to the block to be processed. In the case of HIPI, the logic cannot be added, because of the non-availability of HIPI APIs to know corresponding adjacent block information. The results are depicted in Figure 8, which compare the performance of XHAMI with customized Hadoop MR and HIPI. The results indicate that the performance of XHAMI is much better, and which is nearly half of the time taken by Customized Hadoop MR, and it is extremely better over HIPI. Customized Hadoop MR is implemented with standard Java Image Processing Library, with few customized features over default Hadoop, like retrieving the adjacent blocks information in Map functions for processing the edge pixels of the blocks. This customization requires additional overheads, increasing both the programming and computational complexities. The additional overheads are mainly due to the transfer of whole data block, which is located in different data nodes, than the one where Map function is to be processed.

HIPI has in built Java Processing APIs for processing the jpg and png image formats. For HIPI, the data are organized as a single large block equivalent to the size of the image, as there is no functionality readily available for retrieving the adjacent blocks information. Because of this reason, the experiments for the larger image size data sets starting from serial numbers 3 and above mentioned in Table VIII could not be conducted. XHAMI overcomes the limitations of both Hadoop, and HIPI, by extending the Hadoop HDFS and MR functionalities with the overlapped data partitioned approach, and MR processing using high-level APIs with the integrated open-source GDAL package for handling several types of image formats. XHAMI not only simplifies the programming

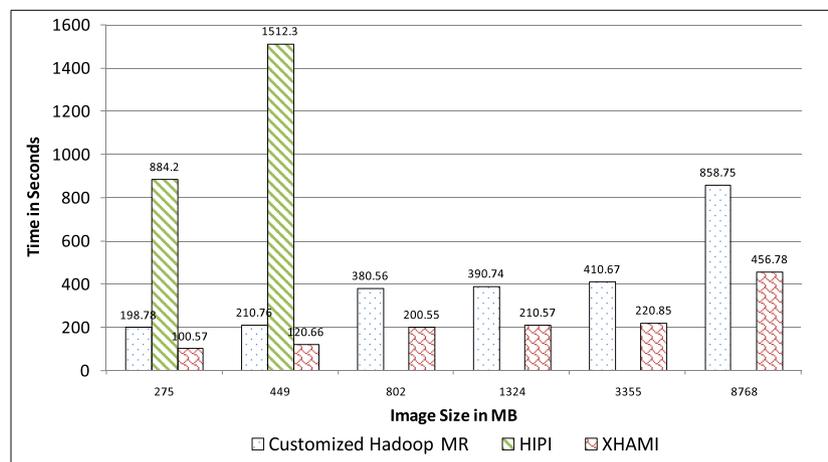


Figure 8. Sobel filter performance.

Table VIII. Read/write performance overheads.

Serial no.	Image size (MB)	Write (s)		Read (s)	
		Default Hadoop	XHAMI	Default Hadoop	XHAMI
1	275	5.865	5.958	10.86	10.92
2	449	14.301	14.365	19.32	19.45
3	802	30.417	30.502	40.2	40.28
4	1324	44.406	77.153	50.28	50.95
5	3355	81.353	88.867	90.3	90.6
6	8768	520.172	693.268	550.14	551.6

complexity but also allows the development of image processing applications quickly over Hadoop framework using HDFS and MR.

4.2. Read/write overheads

Performance of read and write function in default Hadoop and XHAMI with overlap of five scan lines in horizontal partition direction is shown in Figures 9 and 10, respectively. The results shown in Figure 10 represent a negligible read overhead, as the scan lines to be skipped are very few, and




Figure 9. Image write performance.




Figure 10. Image read performance.

also, the position of those lines to skip is known prior. Write performance overheads are shown in Figure 9 and indicate that XHAMI has minimal overheads compared with default Hadoop, and it is observed that data sets in serial nos. 1, 2, 3, and 5 are less than 5%, and for other data sets, it is 33%.

The data sets 4 and 6 are with larger scan line lengths; this in turn has consumed more storage disks space, resulting more read and writes overheads. Performing the vertical direction partition also has resulted in more write overheads for these two data sets, as the number of pixels in vertical direction is more compared with the scan direction. Read performance for all the data sets is less than 0.2%, which is very negligible. Hence, the better mechanism to choose the partitioning approach is used to compute the number of blocks during data partition either in horizontal or in vertical direction and subsequently compute the storage overhead for each block followed by all the blocks. Based on the storage overheads, a data partitioning approach selection can be made, for the one resulted in minimal write overheads. But it is to be observed that for image processing explorations, in general, the images are written once and read several times; hence, the writing overhead is one time activity, which is negligible, while comparing with the overall performance achieved while processing.

5. CONCLUSIONS AND FUTURE WORK

Image processing applications deal with processing of pixels in parallel, for which Hadoop and MR can be effectively used to obtain higher throughputs. However, many of the algorithms in image processing, and other scientific computing, require use of neighborhood data, for which the existing methods of data organization and processing are not suitable. We presented an extended HDFS and MR interface, called XHAMI, for image processing applications. XHAMI offers extended library of HDFS and MR to process the single large-scale images with high-level of abstraction over writing and reading the images. APIs are offered for all the basic forms Read/Write and Query of images. Several experiments are conducted on sample of six data sets with a single large size image varying from approximately 288 MB to 9.1 GB.

Several experiments are conducted for reading and writing the images with and without overlap using XHAMI. The experimental results are compared with the Hadoop system, and HIPI shows that although the proposed methodology incurs marginal read and write overheads, due to overlapping of data, however, the performance has scaled linearly and also programming complexity is reduced significantly. Currently, XHAMI has functions for the data partitioning in horizontal direction; it needs to be extended for both in vertical direction and bidirectional (both horizontal and vertical).

The system is implemented with both the fixed length record and the customized split function, which hides the low-level details of handling and processing, spawning more map functions for processing. However, challenges involved in organizing the sequence of executed map functions for aggregations need to be addressed. We plan to implement the bidirectional split also in the proposed system, which would be the requirement for large-scale canvas images. The proposed MR APIs could be extended for many more image processing and computer vision modules. It is also proposed to extend the same to multiple image formats in the native format itself.

Currently, image files are transferred one at a time from the local storage to Hadoop cluster. In future, data-aware scheduling discussed in our earlier work 29 will be integrated for the large-scale data transfers from the replicated remote storage repositories and performing group scheduling on the Hadoop cluster.

ACKNOWLEDGEMENTS

We express our thanks to Mr. Nooka Ravi of ADRIN for discussion on GDAL library and Mr. Rama Krishna Reddy V. for discussions on image filters based on spatial and frequency domains. We thank Smt. GeetaVaradan, Ex-Director ADRIN, Sri. Santanu Chowdhury, Director ADRIN, and Dr. S. VenkatRaman, Group Director ADRIN, for their support and encouragement in pursuing the research. This paper is a substantially extended version of our nine-page conference paper 30.

REFERENCES

1. IDC, The Digital Universe in 2020: Big Data, bigger digital shadows, and biggest growth in the Far East, Available from: <http://www.emc.com/leadership/digital-universe/index.htm> [last accessed January 2016].
2. Kune R, Konugurthi PK, Agarwal A, Chillarige R, Buyya R. The anatomy of Big Data computing. *Software: Practice and Experience (SPE)* 2016; **46**(1): 79–105 Wiley Press, New York, USA, Jan. 2016.
3. Bakshi K. Considerations for Big Data: architecture and approach, Proceedings of the 2012 IEEE Aerospace Conference – Big Sky, Montana, USA, March 2012.
4. Michael K, Miller KW. Big Data: new opportunities and new challenges. *IEEE Computer* 2013; **46**(6): 22–24.
5. The Apache Hadoop Project. Available from: <http://hadoop.apache.org> [last accessed January 2016].
6. Kelly J. Big Data: Hadoop, business analytics and beyond, Wikibon White paper, 27th August 2012. Available from: <http://wikibon.org/wiki/v/BigData:HadoopBusinessAnalyticsandBeyond> [last accessed January 2016].
7. Ghemawat S, Gobioff H, Leung ST. The Google file system, In Proc. 9th ACM Symposium on Operating System Principles (SOSP 2003), NY, USA 2003.
8. Schwachko K, Kuang H, Radia S, Chansler R. The Hadoop distributed file system, Proceedings of the 26th IEEE Symposium on Mass Storage Systems and Technologies (MSST 2010), Incline Village, Nevada, USA, May 2010.
9. Dean J, Ghemawat S. MapReduce: simplified data processing on large cluster. *Communications of the ACM* 2008; **51**(1): 107–113.
10. Jin C, Buyya R. MapReduce Programming Model for NET-based Cloud Computing, Euro-Par 2009 Parallel Processing, Lecture Notes in Computer Science, 5704 (2009): 417–428.
11. Ekanayake J, Pallickara S, Fox G. MapReduce for data intensive scientific analyses, Proc. IEEE fourth International Conference on e-Science, Indiana Police, Indiana, USA, Dec 2008.
12. Satellite imaging corporation. Available from: <http://www.satimagingcorp.com/satellite-sensors> [last accessed February 2016].
13. Ma Y, Wu H, Wang L, Huang B, Ranjan R, Zomaya A, Jie W. Remote sensing Big Data computing: challenges and opportunities. *Future Generation Computer Systems* 2015; **51**: 47–60.
14. Gonzalez RC, Woods RE. *Digital Image Processing*, 3rd edn, (chapters 1, and 3). Pearson Prentice Hall: Upper Saddle River, NJ, 2007.
15. Cao X, Wang S. Research about image mining technique. *Journal of Communications in Computer and Information Sciences* 2012; **288**: 127–134.
16. Sweeney C, Liu L, Arietta S, Lawrence J. HIPI: a Hadoop image processing interface for image-based MapReduce tasks, undergraduate thesis, University of Virginia, USA.
17. Ekanayake J, Li H, Zhang B, Gunarathne T, Bae S, Qiu J, Fox G. Twister: a runtime for iterative MapReduce, Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing (HPDC 2010), Chicago, Illinois, USA 2010.
18. Jiang W, Ravi VT, Agarwal G. A Map-Reduce system with an alternate API for multi-core environments, Proceedings of the 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing (CCGRID 2010), Melbourne, May 2010.
19. Kennedy L, Slaney M, Weinberger K. Reliable tags using image similarity: mining specificity and expertise from large-scale multimedia databases, Proceedings of the 1st Workshop on Web-Scale Multimedia Corpus, Beijing, October 2009.
20. Shi LL, Wu B, B Wang, XG Yang. Map/Reduce in CBIR applications, Proceedings of the International Conference on Computer Science and Network Technology (ICCSNT), Harbin, December 2011.
21. Yang CT, Chen LT, Chou WL, Wang KC. Implementation on cloud computing, Proceedings of the IEEE International Conference on Cloud Computing and Intelligence Systems (CCIS 2011), Beijing, September 2011.
22. Kocalkulak H, Temizel TT. A Hadoop solution for ballistic image analysis and recognition, Proceedings of the International Conference on High Performance Computing and Simulation (HPCS 2011), Istanbul, July 2011.
23. Almeer MH. Cloud Hadoop MapReduce for remote sensing image analysis. *Journal of Emerging Trends in Computing and Information Sciences* 2012; **3**: 637–644.
24. Demir I, Sayar A. Hadoop optimization for massive image processing: case study of face detection. *International Journal of Computers and Communications and Control* 2014; **9**(6): 664–671.
25. White T. The small files problem: 2009. Available from: <http://www.cloudera.com/blog/2009/02/02/the-small-files-problem> [last accessed February 2016].
26. Potisepp K. Large scale image processing using MapReduce. MSc. Thesis, Institute of Computer Science, Tartu University, 2013.
27. GDAL, Geospatial data abstraction library. Available from: <http://www.gdal.org/> [last accessed January 2016].
28. Apache HBASE project. Available from: <http://www.hbase.apache.org> [last accessed January 2016].
29. Kune R, Konugurthi P, Agarwal A, Rao C, Buyya R. Genetic algorithm based data-aware group scheduling for Big Data clouds, Proceedings of the International Symposium on Big Data Computing (BDC 2014), London, December 2014.
30. Kune R, Konugurthi P, Agarwal A, Rao CR, Buyya R. XHAMI – extended HDFS and MapReduce interface for image processing applications. Proceedings of the 4th International Conference on Cloud Computing for Emerging Markets (CEEM 2015, IEEE Press, USA), Bangalore, India, November 25–27, 2015.