



DISTSENSING: A NEW PLATFORM FOR TIME SERIES PROCESSING IN A DISTRIBUTED COMPUTING ENVIRONMENT

DistSensing: Uma nova Plataforma para o Processamento de Séries Temporais em um Ambiente Computacional Distribuído

**Sávio S. Teles de Oliveira¹, Marcelo de C. Cardoso¹,
Wisllay M. V. dos Santos¹, Paulo C. P. Costa¹,
Vagner J. do Sacramento Rodrigues¹ & Wellington S. Martins²**

¹GoGeo

Rua Leopoldo Bulhões, esquina com a Rua 1014

Quadra 31, Lote 07, Sala 9 Setor Pedro Ludovico - CEP 74820-270 - Goiânia - GO - Brazil
{marcelo.cardoso, vagner, savio.teles, paulo.cezar}@gogeo.io, wisllay.vitrio@gmail.com

²Universidade Federal de Goiás – UFG

Instituto de Informática

Alameda Palmeiras, Quadra D, Câmpus Samambaia - 131 - CEP 74001-970 - Goiânia - GO - Brazil
wellington@inf.ufg.br

*Received on February 20, 2017/ Accepted on June 3, 2017
Recebido em 20 de Fevereiro, 2017/ Aceito em 3 de Junho, 2017*

ABSTRACT

Time series analysis of remote sensing images are indispensable in identifying patterns, trends and changes, and allows the modeling and prediction of events on Earth's surface. For applications with large volumes of data, this analysis should be done in an automated way allowing spatio-temporal filtering in the image database. This paper proposes a new platform, DistSensing, to process these analysis using spatial and relational distributed indices. The DistSensing platform had better performance than the solutions found in the literature when it is necessary run queries in the database using temporal and spatial filters.

Keywords: Remote Sensing, Distributed Systems, Time Series Analysis.

RESUMO

Análises da série temporal de imagens de sensoriamento remoto são indispensáveis na identificação de padrões, tendências e mudanças, além de permitir a modelagem e previsão de eventos na superfície terrestre. Para aplicações com grande volumes de dados, esta análise deve ser feita de forma automatizada permitindo realizar filtragens espaço-temporais na base de dados de imagens. Este trabalho propõe uma nova plataforma, DistSensing, para processar estas análises utilizando índices espaciais e relacionais distribuídos. A plataforma DistSensing teve desempenho melhor que as soluções encontradas na literatura quando existe a necessidade de executar consultas na base de dados utilizando filtros temporais e espaciais.

Palavras-chave: Sensoriamento Remoto, Sistemas Distribuídos, Análise de Séries Temporais.

1. INTRODUCTION

Earth's surface is changing at an unprecedented rate, with alarming disappearance of much of the forest ecosystem, while urban and farming areas are expanding around natural spaces. A time series analysis of remote sensing images is essential to detect these changes (NEVES *et al.* 2015). This entails, for example, providing information on shifts in the spatial distribution of bio-climatic zones, and indicating the variations in large-scale global circulation patterns or changes in land-use.

It is essential to make users aware of both the spatial and temporal dimensions in a Geographic Information System (GIS), because these may reveal implicit relationships which match the reality of the analyzed data (DE OLIVEIRA *et al.*, 2012). The GIS is a computer-based tool for mapping and analyzing feature events on earth, which must allow users to conduct a time-series analysis of remote sensing data filtering in specific geographical regions and at regular time intervals, for example, to trace the evolution of deforested areas in the Amazon forest from 1991 to 1997.

The incorporation of latest-generation sensors to airborne and satellite platforms has led to a nearly continuous stream of data (SMITS *et al.*, 2004), and this sharp rise in the amount of collected information has raised new processing challenges with regard to the time-series analysis of remote sensing data. In addressing these computational requirements, several research endeavors have recently been geared to incorporating high performance computing models in the remote sensing field (PLAZA *et al.*, 2007).

Our work introduces a new platform called DistSensing, to perform the distributed processing involved in time series analysis of remote sensing data, which allows users to have real time spatio-temporal filters. DistSensing obtained more impressive performance gains than other time series processing platforms found in the literature - Van Den Bergh *et al.* (2012), Song *et al.* (2015) - when spatiotemporal filters are required.

The main contributions to this field of study made by this work are as follows: i) it sets out a strategy for partitioning the remote sensing images into cluster nodes; ii) it includes a search algorithm for the processing of the time series

analysis of remote sensing images by means of distributed indexing.

The remainder of this work is structured as follows. Section 2 describes strategies found in the literature to process remote sensing images. Section 3 provides a review of time series processing based on remote sensing data. Section 4 outlines the DistSensing platform. Section 5 describes the methodology and the experimental results. Section 6 summarizes the conclusions and includes a brief description of further work that will be carried out in the future.

This article is based on Oliveira *et al.* (2016), previously presented at GEOINFO conference (<http://www.geoinfo.info/geoinfo2016>).

2. RELATED WORK

Ferreira *et al.* (2016) created a new RDF vocabulary to access spatiotemporal datasets from different kinds of data sources, including remote sensing images. Some works sought to automatically detect changes by using remote sensing images, like Neagoe *et al.* (2014) who adopted non-supervised approaches with less manual intervention. In (Romani *et al.* 2009) a new algorithm was designed to reveal changes in weather patterns by plotting the pixels location from an image that was created to partition them. All these works relied on a single server to process the images, instead of using the computing resources in a cluster.

Several studies employed the MapReduce model. Da Silva Ferreira *et al.* (2015) introduced the architecture of a distributed platform named InterIMAGE Cloud Platform (ICP) to handle with very large volumes of data using clusters of low-cost computers with the Hadoop framework. In LV *et al.* (2010) a K-means clustering algorithm was implemented for remote sensing images and in (Wang *et al.* 2012) a classification algorithm was created for high resolution remote sensing images. The MapReduce model was also used by Almeer (2012) to create a parallel processing platform for remote sensing images based on a cloud computing system. In Lin *et al.* (2013), a platform was proposed and implemented using Hadoop to process remote sensing algorithms with MapReduce models. The work by Rathore *et al.* (2015) relies on an architecture to analyze remote sensing images in real time with Big Data using Hadoop. However, these works fail to

provide any solution for conducting time series analysis by means of remote sensing images.

Song *et al.* (2015) created a Spatiotemporal platform for time series processing of spatial objects, through cloud computing. This includes HDFS, which is used as a distributed file system, and a MapReduce based computing service, which is used to analyze spatial data. This paper does not discuss spatial and temporal indexing techniques nor how the image distribution is carried out in the cluster. When using HDFS, it is important to define the data distribution algorithm, since images from the same geographical region and collected during the same time interval, should be stored in the same machine (VAN DEN BERGH *et al.* 2012). This is to avoid generating a high volume of network traffic while the time series analysis is being processed.

Van Den Bergh *et al.* (2012) established the HiTempo platform to assist research in the area of time series analysis based on remote sensing images. The images were stored as 3D objects, with geographically closed images obtained in the same time interval, being stored in the same machine. HiTempo was designed to make possible the evaluation and comparison of remote sensing algorithms and, for that reason, the timespace filters are only applied before inserting the images into the platform, to prevent the filters from being applied at query time. This means that, it is not possible to perform spatial and temporal filtering after the data insertion and the query cannot apply filters to the image set stored in the HiTempo platform.

3. TIME SERIES ANALYSIS OF REMOTE SENSING IMAGES

There is a wide range of satellites and sensors with different resolutions. In our study, remote sensing images were chosen from Landsat-8, because they are public and easy to access. Each scene taken by a satellite in a given geographical area and time is called a scene. Each scene from Landsat-8 consists of nine different spectral bands, covering an area of around 170 km from north to south and 183 km from east to west.

A time series analysis of remote sensing images collected by satellite is needed to evaluate the features of terrestrial surfaces. A time series is a sequence of images, collected from successive,

and usually uniformly spaced, time intervals. A time series analysis includes methods that are employed to identify patterns and trends, to detect changes, to cluster and model data.

In Figure 1, it is possible to visualize the four dimensions of the time series of remote sensing images. In this graph, the x and y axes represent the spatial limits of each band from a scene, where x is the scene’s geographical longitude, and y is the scene’s geographical latitude. The “bands” axis represents the spectral bands of each scene, and the “time” axis represents the satellite images obtained from various points of time.

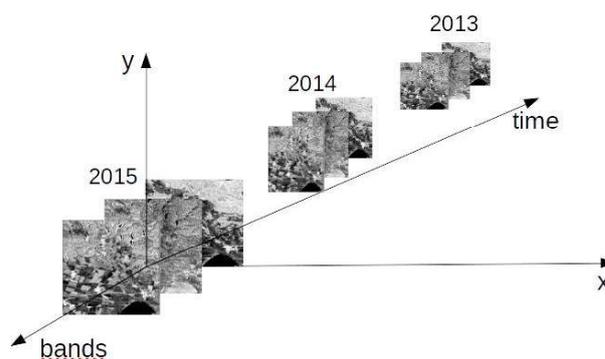


Fig. 1 - Dimensionality of a remote sensing time.

To validate the DistSensing platform in this work, an algorithm for a time series analysis of remote sensing images was implemented the NDVI (Normalized Difference Vegetation Index) historical analysis. The NDVI is a numerical indicator correlated with certain physical properties of the vegetation canopy: leaf area index, vegetation condition, and biomass (CARLSON *et al.* 19970). Through this analysis, it is possible to visualize, for example, the increase in deforestation in certain regions throughout the period by looking at their NDVI time series (MORTON *et al.* 2005). The final result, for instance, is a graph, containing the NDVI mean for each date, as can be seen in the example from Figure 2.

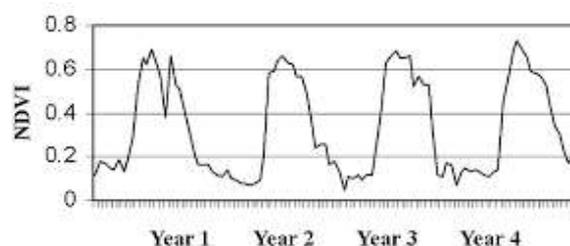


Fig. 2 - Changes in the NDVI in a historical series.

NDVI is derived from spectral reflectance measurements acquired in the visible (RED) and near-infrared (NIR) regions, where NIR and RED are the reflectance in infrared and red spectral intervals, respectively. NDVI is calculated from the difference between RED and NIR bands, normalized by the sum of the same bands. The result of the NDVI ranges between -1 and +1.

4. DISTSENSING: A NEW PLATFORM FOR EFFICIENT DISTRIBUTED PROCESSING OF REMOTE SENSING TIME SERIES

This paper creates a new platform, DistSensing, for processing remote sensing time series in a distributed manner. DistSensing has an highly available, elastic and fault-tolerant architecture (Section 4.1) and provides efficient distributed algorithms for storing, indexing (Section 4.2) and querying images at its databases (Section 4.3).

4.1. DistSensing Architecture

DistSensing has a peer-to-peer architecture in which the cluster servers do not share CPUs, hard drives or memory and all communication is carried out via a message passing system. It is composed by a client layer and a server layer. Client applications interact with the platform through an API for updating and querying the databases. Each client application must use a client library for DistSensing API so that the platform available services can be used. Read and write requests can be sent to any server of the cluster. The list of servers can be obtained from the platform name service.

The server that receives the client request becomes the coordinator for that specific request. The coordinator acts as a proxy between the client application and the cluster servers. To ensure high availability, when the coordinator is unavailable another server is chosen as coordinator. The client library sends the request, receives the response and transfers the result of the operation to the client application.

Cluster nodes exchange information with each other every second using the gossip protocol (SUBRAMANIYAN *et al.* 2006). The Gossip protocol is used for finding out and sharing location and status information about other

cluster servers. Each server exchange messages not only about its status and other factors but also regarding other cluster data nodes with up to three other servers. Thus, all the cluster nodes become aware of the status of the other nodes of the cluster quickly. Failure detection occurs on the basis of the data exchanged via the gossip protocol. The platform uses this information to avoid requesting processing to unavailable servers.

DistSensing stores data replicas on multiple nodes to ensure fault tolerance and reliability. All the replicas are equally important, since there is no primary or master replica. Every object stored is given a single identifier key generated via hashing. Each server is responsible for an equal range of cluster keys, that are accessed through a Distributed Hash Table (DHT) (KARGER *et al.* 1997). When an object is required, the elected server is chosen by means of the DHT.

To ensure platform elasticity, machines might be added or removed from the cluster at any moment. Whenever this happens, the DHT is rebuilt to reflect the new configuration of the cluster. When a new server joins the cluster it queries the DHT and sends a message to the server with the highest disk usage in the cluster to obtain half of the keys kept in this server. Upon removal of the server S its keys are distributed among the remaining servers and its objects are copied from replicas.

4.2. Distribution and Storage of Remote Sensing Images on Cluster Servers

The architecture of DistSensing allows several remote sensing images to be inserted in a distributed and parallel manner. Figure 3 shows step-by-step what happens when an image is inserted in the platform. Insertion starts with a client sending the image bands to the DistSensing platform using the API. Each image band is then sent to a randomly chosen server of the cluster.

In each server, each spectral band of the image is broken into smaller fixed-size blocks that allow a parallel processing of images during the queries. The block size directly affects the size of the job that has to be carried out in the query operations involving the remote sensing images. The greater the number of cores in the cluster servers, the smaller the remote sensing image

block size must be to increase the parallelization and improve the efficiency of query processing. This block size must be set from tests conducted in the cluster environment where the platform is installed. To ensure reliability and fault tolerance, as well as providing load balancing, each block is stored on a server and replicated in two other servers.

As Figure 3 shows, block distribution is based on four dimensions of each block: the first two are x and y, and represent the spatial limits of the block, the third is the time and the fourth one is related to the image band. Each server in the cluster has its own storage system regardless of the others servers, where the blocks are stored.

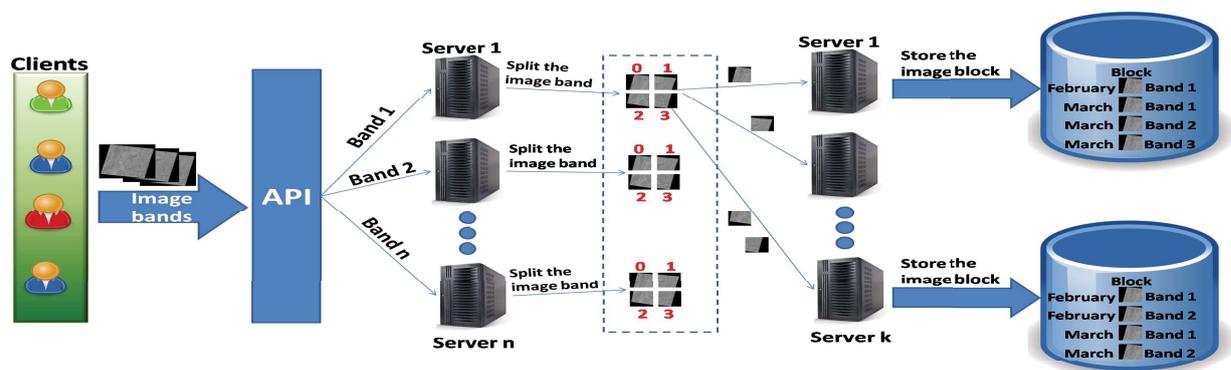


Fig. 3 - Distribution and Storage of Remote Sensing Images on DistSensing Platform.

Thus, blocks from the same geographical area collected at different times are stored at the same server, and this reduces network traffic during the execution of time series analysis. By means of this block distribution, the remote sensing algorithms that carry out the time-series analysis in each image pixel (e.g. NDVI) do not have to send data over the network, since the historical series of each pixel is stored in a server.

Each image block is linked to a geographical location and contains metadata from the original image from which it was generated, such as the date when the image was captured. A Quad Tree (FINKEL *et al.* 1974) spatial index is built in each server of the cluster, allowing spatial filters to be performed. The index in each server is built from each spatial image block bounding stored in that server. The DistSensing platform is responsible for coordinating the distributed spatial filter amongst the several spatial indices.

Quad Tree is a tree data structure in which each internal node has exactly four children and it was chosen because of its ability to filter remote sensing data (FU *et al.* 2013). The key idea is to partition the space into four quadrants. Each node in the tree either has exactly four children, or has no child. Each child represents a subquadrant of its parent. Quadrants that can

no longer be subdivided are represented as leaf nodes that contain data corresponding to that specific subregion.

In addition, an index is built from the metadata (remote sensing image attributes, like the date) of each image in the database. This index is important since it provides with ways to, for instance, conduct an analysis at specific time intervals. The Apache Lucene library was used to build this index. It is available at each server of the cluster and the DistSensing platform is responsible for coordinating the distributed filter amongst the Lucene search engines. The client sends a request that defines the restrictions using Apache Lucene query syntax, which allows it to filter images through phrase, wildcard, range and full-text-search queries.

4.3. Querying the Remote Sensing Database

The query execution process starts with the client sending the request to the DistSensing platform with the spatial filters and metadata attributes constraints. One of the platform servers is then chosen as the Master that coordinates the execution of that request. Figure 4 describes the steps that must be followed inside the DistSensing platform.

The query request is sent to all the cluster

servers, since each server stores remote sensing image blocks. After this, each server locally processes the spatial filters by means of the Quad-Tree and the image metadata, by using the Apache Lucene on the basis of client-specified restrictions. The result is stored and sent to the Master server as a list of ids of image blocks, along with metadata from the blocks necessary for the execution.

The Master server is responsible for receiving the ids and metadata and removing

duplicates, since the data is replicated among the servers. For each resulting pair containing id and metadata, the DistSensing platform queries its DHT to discover which servers are storing the image block with that id. This DHT stores the id of each object as its key and the IP addresses of servers as value. A server S is randomly chosen from this list and combined with that id. For each server S, the list of ids and metadata of the images stored by S is retrieved and then sent to server S.

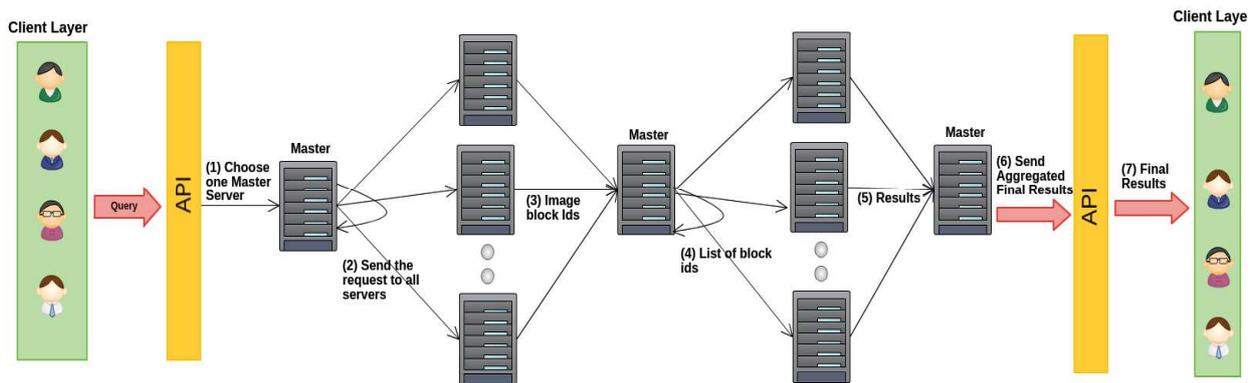


Fig. 4 - Steps to Querying the Remote Sensing Database.

In each server S, the blocks are retrieved from the storage system and the remote sensing time-series analysis is conducted with the algorithm requested by the client, such as the NDVI time-series analysis. The results obtained from processing this block list, are stored and sent to the master server. The master server receives replies from each server and aggregates them in order to create the final result, which is then sent to the client. This aggregation algorithm is defined by the clients of the DistSensing platform in accordance with their requirements.

On the basis of the analysis shown in Section 3, which employs the normalized difference vegetation index (NDVI), the client sends the spatial and metadata restrictions to the DistSensing platform. The date attribute is then used to select remote sensing image blocks in a given time slice. The coordinator sends the client request to each server of the cluster which calculates the average NDVI for each image block. Once the average NDVI is known for all the image blocks, the values are aggregated by date to create the average NDVI for each date.

These aggregate values are sent to the Master server which processes the final aggregation on the basis of data obtained from all the servers and generates the average NDVI for each date. These values are sent to the client which is then capable of creating graphs of NDVI variations over a period of time.

5. PERFORMANCE EVALUATION

An algorithm was implemented to evaluate the performance of the DistSensing platform and thus allow an analysis of the temporal variations in the NDVI values, as shown in Section 3. The response times from the DistSensing platform are compared with those of the HiTempo platform recommended by Van Den Bergh *et al.* (2012), which is similar to the proposal put forward by Song *et al.* (2015), that are unable to perform temporal and spatial filters during the query execution. The image processing module from the solution found by Van Den Bergh *et al.* (2012) was implemented to evaluate a HiTempo platform with spatial and temporal constraints and in the final stage.

5.1. Experiments and Database

A horizontal scalability test was conducted to measure the performance of the DistSensing platform when adding servers to the cluster. Ideally the system should have linear horizontal scalability, which is not always possible: a) since the tasks distributed in the cluster have to be synchronized and b) the devices are subject to physical constraints as in the case of those used for networking. The tests were run on Intel Core i7 3.4 GHz machines, with 16 GB of RAM and 1 TB hard drives. The machines were connected to a 1Gbit/s Ethernet network and a 6248P Dell PowerConnect switch. 53 remote sensing images from Brazil were collected by the Landsat-8 satellite from 2013 to 2015, each image with approximately 2 GB. Each of the image’s band was broken into blocks of 1024 x 1024 pixels during the insertion procedure.

The algorithm for analyzing temporal variation of NDVI values was evaluated with clusters of 1, 2, 4 and 8 servers and took into account several spatial and temporal filter configurations: i) No filter, covering 100% of the database, and mirroring the HiTempo strategy (HiTempo), ii) spatial constraints covering 27,1% of the database (Large_Spatial), iii) spatial constraints covering 6,1% of the database (Small_Spatial), iv) temporal constraints covering 34% of the database (Date_Large), v) temporal constraints covering 5% of the database (Date_Medium), vi) temporal constraints covering 1% of the database (Date_Small), vii) having spatial and temporal constraints that result in an empty response. Names in brackets are used in the charts shown in Section 5.2.

The algorithm was executed 20 times for each test configuration and the response time on the charts is the average of the 20 response times observed. A client library was created so that data could be sent for insertion, and queries performed with spatial and relational constraints which involved measuring the response times for each specific test.

5.2. Evaluation

The first test was carried out to determine the scalability of the DistSensing and HiTempo platforms when processing remote sensing time-series analysis, as shown in Figure 5. The response time was obtained after the time-series analysis

had been conducted on the entire database. Both platforms achieved a lower response time as more machines were added to the cluster. This scalable behavior was made possible by dividing the image into blocks, and allowing a distributed and parallel processing of these blocks on the cluster. The HiTempo platform had a similar behavior and response time, since there was no need for subsequent filtering.



Fig. 5 - Horizontal Scalability.

Figure 6 shows a comparison between HiTempo and DistSensing platform when running the analysis with temporal filters on the image database. Compared with the HiTempo platform, DistSensing achieves response times up to 53 times lower if the “Date_Small” temporal filter is selected for the execution. Since the HiTempo platform does not allow temporal filters at the query time, it must process the complete database and apply the filter later. With “Date_Small” temporal filter, the DistSensing platform was able to filter 99% of the database, leaving only 1% for processing the analysis on remote sensing images.

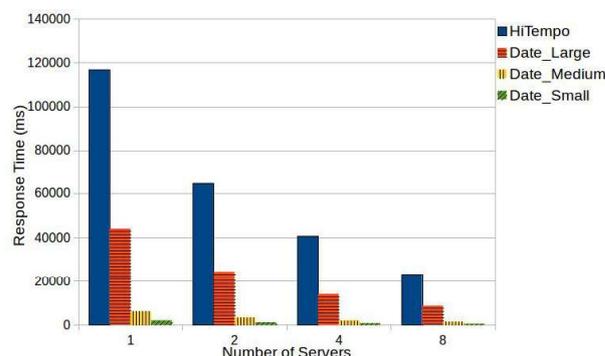


Fig. 6 - Comparison between DistSensing platform and the solution proposed by HiTempo platform.

The DistSensing platform shows scalability when processing the algorithm with temporal filters, especially with filters that are less

restrictive like “Date_Large” where there is a greater need for computational resources during the analysis of the image time series. As shown in Figure 7, the DistSensing platform shows scalability when processing algorithms with more restrictive filters than “Date_Large”, although not at the same levels of scalability. However, adding more servers for restrictive filters like “Date_Small” and “Date_Medium” do not result in proportionally lower response times, unless the image database increases and requires more computational resources.

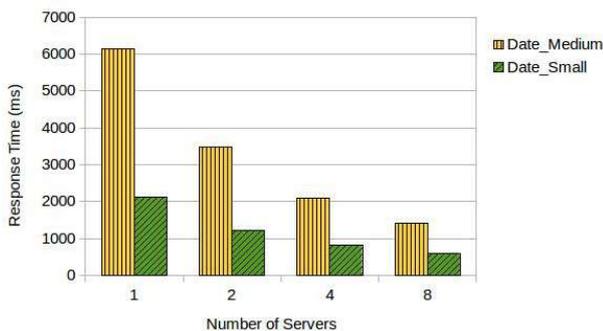


Fig. 7 - Scalability on temporal filters.

Spatial filters are important for remote sensing specialists since they allow them to conduct an analysis that takes into account their geographic context and interest. Figure 8 shows the performance of the DistSensing platform when the algorithm is processed with spatial filters. The DistSensing platform was up to 9 times faster than the solution proposed by the HiTempo platform when the spatial filter “Small_Spatial” was used by the algorithm. With the “Date_Small” temporal filter DistSensing was up to 53 times faster because this filter covers 1% of the database while the “Small_Spatial” spatial filter corresponds to 6,1%.

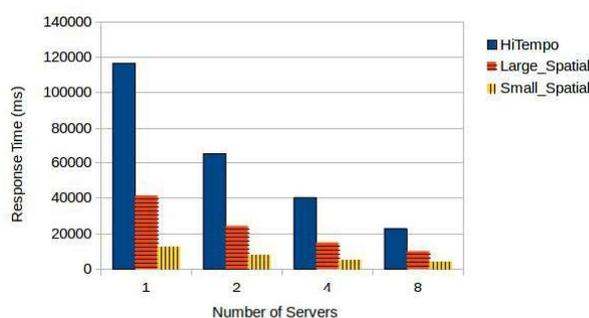


Fig. 8 - Database spatial filtering with *HiTempo* and *DistSensing* platforms.

Tests were carried out with spatial and temporal constraints which lead to empty responses. The DistSensing platform showed constant response times of approximately 10 ms in these cases, regardless of the cluster size. The HiTempo platform performed up to 11,600 times worse, with response times of 116 seconds when only one server was used. With 8 servers HiTempo took 23 seconds to process the data, which is a performance 2300 times worse than that of DistSensing. The reason for this huge difference is that DistSensing is capable of filtering out images that do not meet the constraints, before processing the algorithm on the remote sensing images.

The algorithm employed for analysing the temporal variation of the NDVI values, was found to be scalable, both with regard to the HiTempo platform and the DistSensing platform, as more servers were added to the cluster. The DistSensing platform had a better performance when temporal and spatial filters were used, with a speed up to 53 times faster when a temporal filter was used that corresponded to 1% of the database, and up to 9 times faster if a spatial filter was used that corresponded to 6,1% of the database. In cases where spatial and temporal filters lead to an empty dataset, the DistSensing platform was up to 11,600 times faster.

6. CONCLUSION

Time series analysis is crucial for detecting patterns, trends and changes, as well as providing us with ways to model and predict events on the earth’s surface. The automatic execution of an analysis of remote sensing time series has become a challenge due to the increase in the amount of remote sensing data. To the best of our knowledge, there is no study on the methods employed for processing a time series analysis that allows the filtering of data based on a geographical region and time period.

This paper proposes a new platform called DistSensing, to conduct an analysis of remote sensing time series in a distributed manner. Spatial and relational indices were built to provide query functionality at the image database. DistSensing allows patterns, trends and changes on Earth’s surface to be detected with lower response times. When temporal filters are used, DistSensing is up to 53 times faster than the

HiTempo platform proposed by Van Den Bergh *et al.* (2012), and if spatial filters are needed, the performance can be up to 9 times better than the HiTempo platform. Furthermore, if the combination of spatial and/or temporal filters generates an empty dataset, the DistSensing platform shows a response time that is 11,600 times lower than those of the HiTempo platform.

Tests on bigger clusters and with a larger volume of data will be conducted in future research. In addition, experiments will be carried out to find out what is the ideal size for the image blocks on the basis of the cluster configuration.

REFERENCES

- ALMEER, M. H. Cloud *hadoop* map reduce for remote sensing image analysis. **Journal of Emerging Trends in Computing and Information Sciences**, v. 3, n. 4, p. 637-644, 2012.
- CARLSON, T. N. AND RIPLEY, D. A. On the relation between ndvi, fractional vegetation cover, and leaf area index. **Remote sensing of Environment**, v. 62, n. 3, p. 241-252, 1997.
- DA SILVA FERREIRA, R., OLIVEIRA, D. A. B., HAPP, P. N., DA COSTA, G. A. O. P., FEITOSA, R. Q., AND BENTES, C. Interimage cloud platform: Em direção à arquitetura de uma plataforma distribuída e de código aberto para a interpretação automática de imagens baseada em conhecimento. **XVII Simpósio Brasileiro de Sensoriamento Remoto - SBSR**, p. 5264-5271, 2015.
- DE OLIVEIRA, M. G. AND DE SOUZA BAPTISTA, C. Geostat- A system for visualization, analysis and clustering of distributed spatiotemporal data. In **XIII Brazilian Symposium on Geoinformatics (GEOINFO)**, p. 108-119, 2012.
- FERREIRA, K. R., DE OLIVEIRA, A. G., MONTEIRO, A. M. V. AND DE ALMEIDA, D. B. F. Temporal GIS and spatiotemporal data sources. **Revista Brasileira de Cartografia**, v. 68, n. 6, p. 1191-1202, 2016.
- FINKEL, R. A. AND BENTLEY, J. L. *Quad trees* a data structure for retrieval on composite keys. **Acta informatica**, v. 4, n. 1, p. 1-9, 1974.
- FU, G., ZHAO, H., LI, C., AND SHI, L. (2013). Segmentation for high-resolution optical remote sensing imagery using improved quadtree and region adjacency graph technique. **Remote Sensing**, v. 5, n. 7, p. 3259-3279, 2013.
- KARGER, D., LEHMAN, E., LEIGHTON, T., PANIGRAHY, R., LEVINE, M., AND LEWIN, D. Consistent hashing and random trees: Distributed caching protocols for relieving hot spots on the world wide web. In **Proceedings of the twenty-ninth annual ACM symposium on Theory of computing**, 1997. p. 654-663.
- LIN, F.-C., CHUNG, L.-K., WANG, C.-J., KU, W.-Y., AND CHOU, T.-Y. Storage and processing of massive remote sensing images using a novel cloud computing platform. **GIScience & Remote Sensing**, v. 50, n. 3, p. 322-336, 2013.
- LV, Z., HU, Y., ZHONG, H., WU, J., LI, B., AND ZHAO, H. Parallel k-means clustering of remote sensing images based on *mapreduce*. In **Proceedings of the 2010 International Conference on Web Information Systems and Mining, WISM'10**, p. 162-170, 2010.
- MORTON, D. C., DEFRIES, R. S., SHIMABUKURO, Y. E., ANDERSON, L. O., DEL BON ESPÍRITO-SANTO, F., HANSEN, M., AND CARROLL, M. Rapid assessment of annual deforestation in the Brazilian Amazon using MODIS data. **Earth Interactions**, v. 9, n. 8, p. 1-22, 2005.
- NEAGOE, V., CIUREA, A., BRUZZONE, L., AND BOVOLO, F. A novel neural approach for unsupervised change detection using SOM clustering for pseudo-training set selection followed by CSOM classifier. In **Geoscience and Remote Sensing Symposium (IGARSS), 2014 IEEE International**, p. 1437-1440, 2014.
- NEVES, A. K., BENDINI, H. N., KÖRTING, T. S., AND FONSECA, L. M. G. (2015). Combining time series features and data mining to detect land cover patterns: a case study in northern Mato Grosso state, Brazil. In **XVI Brazilian Symposium on Geoinformatics (GEOINFO)**, pages 174-185.
- OLIVEIRA, S.; CARDOSO, M.; SANTOS, W.; COSTA, P.; SACRAMENTO, V.; MARTINS, W. A new platform for time-series analysis of remote sensing images in a distributed computing

- environment. In: **XVII Brazilian Symposium on Geoinformatics (GEOINFO)**, p. 128–139, 2016.
- PLAZA, A. J. AND CHANG, C.-I. High Performance Computing in Remote Sensing. **Chapman & Hall/CRC**, p. 1-496, 2007.
- RATHORE, Muhammad Mazhar Ullah et al. Real-time big data analytical architecture for remote sensing application. **IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing**, v. 8, n. 10, p. 4610-4621, 2015.
- ROMANI, L. A., DE ÁVILA, A. M. H., ZULLO JR, J., TRAINA JR, C., AND TRAINA, A. J. (2009). Mining climate and remote sensing time series to discover the most relevant climate patterns. In **SBBD**. 2009, pages 181–195.
- SMITS, P. AND BRUZZONE, L. Analysis of multi-temporal remote sensing images. **World Scientific**, volume 3, 2004.
- SONG, W., JIN, B., LI, S., WEI, X., LI, D., AND HU, F. Building spatiotemporal cloud platform for supporting gis application. **ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences**, v. 2, n. 4, p. 55, 2015.
- SUBRAMANIYAN, R., RAMAN, P., GEORGE, A. D., AND RADLINSKI, M. Gems: Gossip-enabled monitoring service for scalable heterogeneous distributed systems. **Cluster Computing**, v. 9, n. 1, p. 101-120, 2006.
- VAN DEN BERGH, F., WESSELS, K. J., MITEFF, S., VAN ZYL, T. L., GAZENDAM, A. D., AND BACHOO, A. K. *Hitempo*: a platform for time-series analysis of remote-sensing satellite data in a high-performance computing environment. **International journal of remote sensing**, v. 33, n. 15, p. 4720-4740, 2012.
- WANG, G., HE, G., AND LIU, J. A new classification method for high spatial resolution remote sensing image based on mapping mechanism. In **Proceedings of the International Conference on Geographic Object-Based Image Analysis (GEOBIA' 12)**, pages 186–190, 2012.