

Table of Contents

1.	Performance Issues, and measurements and analysis	2
1.1.	General Performance Issues in Interoperable Service-oriented GIS.....	3
1.1.1.	Using Semi-structured Data Model.....	3
1.1.2.	Data Characteristics	4
1.2.	Ordinary GIS System Performance Using Common Data Model for the Interoperability (To compare with)	5
1.3.	High Performance Design and Evaluation of the System.....	9
1.3.1.	Data-oriented Approaches.....	9
1.3.1.1.	Streaming Data Transfer	10
1.3.1.2.	Pull Parsing and Application Specific Rendering.....	13
1.3.2.	Federator (AWMS)-oriented Approaches.....	16
1.3.2.1.	Pre-Fetching	17
1.3.2.1.1.	Fetching module (PM).....	19
1.3.2.1.2.	Performance Evaluation	21
1.3.2.2.	Locality-based Query Decomposition and Parallel Processing	24
1.3.2.2.1.	Caching, Cached-data Extraction and Rectangulation	26
1.3.2.2.2.	Parallel-Processing with caching	30
1.3.2.2.3.	Performance Evaluation	36

1. Performance Issues, and measurements and analysis

This chapter presents the common performance issues in service-oriented, federated and interoperable GIS systems built based on the common structured data model. As the common data model, OGC defined Geographic Markup Language (GML) [GML] is used. Developing a federated Information system inspired us enhancing the whole system performance by applying novel parallel processing and caching techniques applied together in large scale interoperable information systems (see Chapter 1.3.2). In addition to this, we proposed some other innovative performance enhancement techniques (see Chapter 1.3.1) such as streaming data transfers, and enhanced parsing and rendering of semi-structured geo-data sets (GML). At the end of each chapter explaining these techniques, there are performance tests and analysis.

The organization of the rest of the chapter is as follows. Chapter 1.1 summarizes and reviews the general performance issues of interoperable service-oriented GIS systems in which interoperability is granted by using XML-structured common data model and Web Services. Chapter 1.2 presents the limits of the ordinary GIS systems without having any performance enhancements which will be our comparison base for our proposed techniques. Throughout the document, with the term “ordinary” we mean on-demand, single-threaded and no-caching systems. The last chapter (Chapter 1.3) explains our approaches to developing high performance GIS systems, and performance evaluations by comparing with the ordinary systems. We approach the performance issues from the two aspects. One is data-oriented and the other is federator-oriented. The data-oriented approaches deal with transferring large-sized XML structured data in common model, and high performance parsing and rendering algorithms. The federator-oriented approaches deal with the performance enhancement techniques based on data characteristics. For the un-frequently changing archived data handling we propose pre-fetching

technique. On the other hand, for the frequently changing archived data, we propose a hybrid technique composed of caching and parallel processing applied together. This hybrid system is actually proposes a novel locality based workload forecasting and parallel processing through range query decomposition (Chapter 1.3.2.2) for variable sized and un-evenly distributed data.

1.1. General Performance Issues in Interoperable Service-oriented GIS

Performance issues in interoperable service-oriented GIS can be generalized into two groups:

- Issues regarding semi-structured data model (GML) (Chapter 1.1.1).
- Issues regarding domain specific data characteristics (Chapter 1.1.2). For example in GIS, the location attribute (defined as (x,y) coordinate) is the major attribute of geo-data. Based on the location, the data is described as un-evenly distributed and variable sized.

1.1.1. Using Semi-structured Data Model

Using semi-structured data model enables interoperability and inter-service communication. XML's emergence as the de facto standard for encoding tree-oriented, semi-structured data has brought significant interoperability and standardization benefits to distributed computing. On the other hand, performance has been still a persistent concern for large scale applications, because of the size issues and processing overheads [Lu2006]. The processing is detailed as parsing and differentiating (separating) the core-data from the attributes and other tags to create required application specific data formats.

Structured data representations enable adding some attributes and additional information to the data. These attributes and additions are mostly due to the interoperability and security reasons. XML representations of data tend to be significantly larger than binary representations of the

same data. The larger document size means that the greater bandwidth is required to transfer of data, as compared to the equivalent binary representations. The larger size often implies greater processing costs as well, since much of the overhead involved in communication processing is going to be based on the data volume.

There are two well-known and commonly-used paradigms for processing XML data, the Document Object Model (DOM) and the Simple API for XML (SAX). DOM builds a complete object representation of the XML document in memory. This can be memory intensive for large documents, and entails making at least two passes through the data. SAX operates at one level lower. Rather than actually constructing a model in memory, it informs the application of elements through callbacks. This also requires at least two passes through the data. These are all expensive and resource (such as CPU and memory) consuming processes and they don't provide enough performance for the large scale applications.

1.1.2. Data Characteristics

The different domains have different data types having different characteristic to be handled. As an example, in GIS domain science applications need to manipulate geo-data. Geo-data is described with its location ((x,y) coordinates) on the earth. Based on the location attribute, geo-data is un-evenly distributed (such as human-population and temperature distributions) and variable sized. Because of these characteristics, it is not easy to implement some well-known performance enhancing techniques as applied in other science domains. Since it is not possible to know the work-load earlier, the classic load balancing algorithms do not work for the variable sized and unevenly distributed data. The work is decomposed into independent work pieces, and the work pieces are of highly variable sized. This issue is illustrated in Figure 1 for the case of using one-step-binary query partitioning based on the location attribute of the data. As it is

illustrated in the figure, there are four worker nodes, and the worker node assigned to R2 gets the heaviest part of the total work, and therefore the expected performance gain from using classic load balancing will not be obtained.

The geo-data is queried based on their attributes. Since all the data is described by their locations, in order to get the data sets falling in a specific region, the bounding box (bbox) values are used. The regions are defined in bboxes. A bbox defines a rectangular shape in a two-dimensional coordinate plane, and it is formulated as $(\min_x, \min_y, \max_x, \max_y)$. For example, Figure 1 shows a region formulated in bbox value (a, b, c, d) .




Figure 1: Un-balanced load sharing. Server assigned R2: “ $(\frac{a+b}{2}, \frac{b+d}{2}), (c, d)$ ” gets the most of the work.

1.2. Ordinary GIS System Performance Using Common Data Model for the Interoperability (To compare with)

In order to solve data and service heterogeneities for the GIS computation and data services OGC and ISO/TC-211 standards are used. These standards recommend using structured common data model called GML for the representation of location based geo-data. The standard bodies aim is to make the geographic information and services neutral and available across any network,

application, or platform. Currently the two major geospatial standards organizations are the Open Geospatial Consortium (OGC) and the Technical Committee tasked by the International Standards Organization (ISO/TC211).


With the ordinary system we mean a GIS developed without using any novel advanced techniques to handle the data. Most of the implementations are based on single-threaded and on-demand processing. In order to compare and contrast our novel approaches to creating large scale high-performance GIS system with the ordinary systems, we tested and presented their performance results at Table 1 and *Figure 2*.

This performance results teach us valuable lessons in terms of the capabilities and limits of the general distributed and interoperable GIS systems. From the figure we draw following conclusions. First, for small data payloads (less than 1MB) the response time is acceptable. However for larger data payloads the performance gets worse and the response time gets relatively longer.

Table 1: The round-trip times (or response times) of the ordinary system.

Data Size -MB	Response times in minutes
1	0.97
20	19.25
40	38.17
60	58.20
80	76.75
100	95.80

(a)



(b)




Figure 2: (a) Performance result of the ordinary system. (b) Sample output consisting of two layers.

The ordinary system test above shows that even if we apply advanced parsing and streaming data transfer techniques (Chapter 1.3.1), the performance is still not enough in order to meet Geo-Science Grids' performance requirements. As you see, if the spatial data is over 1MB, ordinary system framework is not feasible to use in large scale science applications. Time column (y) in the *Figure 2* (a) includes querying, transforming, rendering and displaying spatial data. In other words the figure illustrates the response time of the ordinary GIS systems as formulated below:

$$time_{(measured)} = time_{(result\ is\ displayed)} - time_{(client\ makes\ request)}$$

Measured time in the figure ($time_{(measured)}$) can be detailed as below:

- [***time***_(client makes request)] Client makes requests by interactive smart map tools (Geo-Science Portal) for Aggregator Web Map Server (AWMS).
- AWMS parse and render requests and define set of actions required based on the requests and its capabilities file.
- AWMS Creates map images (from the returned datasets) and returns them to the clients:
 - Defines the set of WFSs [WFS] and other WMSs [WMS] to communicate with to build the response by in accordance with its capability file.
 - Creates requests for WFSs and other WMSs
 - Invokes WFSs *getFeature* Web Services for vector data encoded in Geographic Markup language (GML) [GML].
 - Invokes other WMSs *getMap* Web Services for raster data rendered in map images
 - Transferring GML data (feature collections) from WFS and WMS
 - Parsing and rendering returned GML data sets
 - Aggregating and overlaying layers according to the request and capability file.
 - Sending the map images to the WMS Client.
- [***time***_(map is displayed)] Client shows the returned maps on his browser

From our experience we saw that depending on the total data size, over %90 of the *time*_(measured) comes from the step called “transferring GML data (feature collections) from WFS and WMS”. Because of that, even if we use the most efficient and fast parsing and rendering algorithms (such as using pull parsing or application specific XPath querying), it won’t improve performance very much if the data transfer time still stays that much high as shown in the *Figure 2*.

The next chapter presents our novel approaches to enhance the ordinary GIS system’s performances presented here. Our approaches are summarized as data-oriented and federator-

oriented. At the end of every section, we will make evaluations of the proposed techniques by comparing with the ordinary system performance shown in *Figure 2*.

1.3. High Performance Design and Evaluation of the System

Our approaches to the performance issues are grouped into two. First group of approaches deals with the general performance issues result from using semi-structured data encodings (such as GML), and large size data exchange, parsing and rendering (Chapter 1.3.1). Second group of approaches is regarding the federator oriented design and techniques to enhance the overall system performance (Chapter 1.3.2).

1.3.1. Data-oriented Approaches

Distributed GIS systems typically handle a large volume of datasets. Therefore the transmission, processing and visualization/rendering techniques used need to be responsive to provide quick, interactive feedback. There are some characteristics of GIS services and data that make it difficult to design distributed GIS with satisfactory performance. One of them is that GIS services often transmit large resulting datasets such as structured data, images, or large files in tabular-matrix formats.

In order to provide interoperability and extensibility we use common data format represented and formulated in XML. This degrades the performance even worse for large scale applications. The major hurdle of the proposed federated GIS framework is encoding, transferring and rendering the data in common data model. In the following two sub-sections we present our approaches to these issues.

Two approaches are presented here. One is regarding large scale structured data transfer (Chapter 1.3.1.1) and other is regarding the large scale data parsing (Chapter 1.3.1.2).

1.3.1.1. Streaming Data Transfer

Our experience shows that although we can easily integrate several GIS services into complex tasks by using Web Services, providing high-rate transportation capabilities for large amounts of data remains a problem because the pure Web Services implementations rely on SOAP [Donbox] messages exchanged over HTTP. This conclusion has led us to an investigation of topic-based publish-subscribe messaging systems for exchanging SOAP messages and data payload between Web Services. We have used NaradaBrokering [Pallickara2003] which provides several useful features besides streaming data transport such as reliable delivery, ability to choose alternate transport protocols, security and recovery from network failures.

Naradabrokering is a message oriented middleware (MoM) [Tran] system which facilitates communications between entities through the exchange of messages. This also allows us to receive individual results and publish them to the messaging substrate instead of waiting for whole result set to be returned.

In case of transferring the GML result set in the form of string causes some problems when the GML is larger than some amount of size (1 MB see *Figure 2-a*). Since the WFS returns the resulting XML document as an <xsd:string>, this has to be constructed in memory and the size will depend on several parameters such as the system configuration and memory allocated to the Java Virtual Machine etc. Consequently there will be a limit on the size of the returned XML documents. For these reasons we have investigated alternative ways for data transport and, researched the use of topic based publish-subscribe messaging systems for streaming the data. Our research on NaradaBrokering shows that it can be used to stream large amount of data between nodes without significant overhead. Additional capabilities such as reliable messaging and support for different transport protocols already inherent in NaradaBrokering show that it is

a powerful yet easy to integrate messaging infrastructure. For these reasons we have developed a novel Web Map Service and Web Feature Service that integrate OGC specifications with Web Service-SOAP [Donbox] calls and NaradaBrokering messaging system. Architecture is shown in Figure 3.

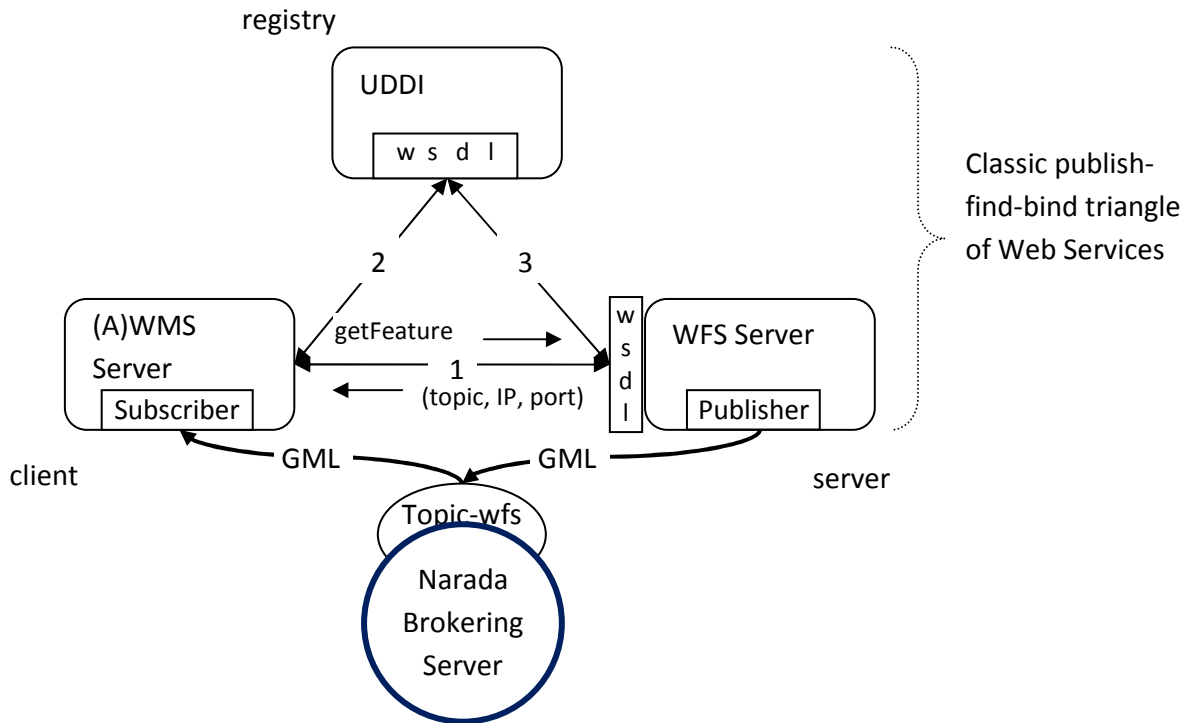


Figure 3: Streaming data transfer using Naradabrokering publish-subscribe topic based messaging middleware.

Connection lines 2 and 3, and UDDI (Universal Description, Discovery and Integration) [uddi] service are displayed in the figure for showing classic publish-find-bind triangle of the Web Service based Service Oriented architecture. We don't go into details of these interactions and UDDI registry service in this document but these can be summarized as following. WFS services publish their existence and service providing with their WSDL service description files (line-3). Clients (such as WMS) find appropriate WFS by searching UDDI registries (line-2). After finding appropriate service, clients are bind to that service by creating their client stubs. Instead

of using lines 2 and 3, clients can also directly communicate with the services if they know the service's WSDL file earlier.

In case of streaming through Naradabrokering, the clients make the requests with standard SOAP messages (line-1) but for retrieving the results a NaradaBrokering subscriber class is used. Through first request to Web Service (called *getFeature*), WMS gets the topic (publish-subscribe for a specific data), IP and port to which WFS streams requested data. Second request is done by *NaradaBrokering* Subscriber. Even in the case of that the whole data is not received by WMS; WMS can draw the map image with the returned data. This depends on the WMS's internal implementation.

Table 2 gives a comparison of the streaming and non-streaming data access approaches for the different sizes of data. These values are obtained by running Pattern Informatics (PI) geo-science application over the earthquake seismic data records. These are data access times including query conversion at WFS, result set conversion from database to GML and transfer times from WFS to AWMS.

Table 2: Data access times while using (1) streaming and (2)non-streaming data transfer techniques.

Data Size (KB)	Time for Streaming	Time for Non-Streaming
10	31.3	3912.5
30	100	3917.1
100	320.1	4098.7
300	826.7	4414
1000	2414.2	5662.6

Figure 2 below explains that streaming data transfer enhancement is still not enough for providing satisfactory large scale application performance. See Chapter 1.3.2 for the proposed overall performance enhancement techniques.




Figure 4: Comparisons of Streaming vs. Non-Streaming data transfer timings.

1.3.1.2. Pull Parsing and Application Specific Rendering

Proposed system includes data rendering/filtering tasks assigned to Web-based Map Services to create comprehensible data representations derived from semi-structured common data (GML). These comprehensible representations are called maps. Regarding the rendering of large GML data and creating map images we use parsers.

There are three general parsing techniques proposed for processing XML structured data. These are document model, push model and pull model. There are also other hybrid alternatives built on these main approaches. In order to process data in XML structured common data model we use pull parsing technique.

Pull parsing, as exemplified by the XML Pull Parser [Alexander], is an efficient paradigm similar to SAX in that it does not build a complete object model in memory. It differs in that the tags and content are returned directly to the application from calls to the parser, rather than indirectly in the form of callbacks. The pull approach of this parsing model results in a very small memory footprint (no document state maintenance required – compared to DOM), and very fast processing (fewer unnecessary event callbacks - compared to SAX).

Pull parser only parses what is asked for by the application rather than passing all events up to the client application as SAX parsing does. You can see the article where pull parsing is compared with other leading Java based XML parsing implementations [Sosnoski].

Pull parsing does not provide any support for validation. This is the main reason that it is faster than its competitors. Since all our services are OGC compatible and created in Web Service principles, we do not necessarily need validation. In OGC, services describe themselves by capability document and servers know each other by exchanging these document. If you are sure that data is valid (as in our case), or if the validation errors are not catastrophic to your system, or you can trust validity of the capabilities document of the server you are in contact, then using XML Pull Parsing gives the highest performance results. For example in communication between WFS and WMS, since we know that WFS provides feature data in OGC's GML format [GML], it is very advantageous skipping validation and using "pull parsing".

For application specific comparison of Pull parsing and DOM see Table 3 and Figure 5. The performance values are measured in milliseconds and data sizes are in KBs. Performance test is done with 1GB allocated JAVA Virtual Machine. Infinite values in the table represent memory exceptions thrown. The figure illustrates the timing values for the data size till 10MB. Above this threshold value DOM become useless for the Virtual Machine allocated 1GB memory.

Pull parsing column of the table does not show parsing values. Inc case of pull parsing, geometry data is extracted from GML without making parsing. For the GML data we use earthquake seismic data records kept in archived stores and provided by WFS deployed at Community Grids Labs in Indiana University.

Table 3: DOM and Pull parsing (Xpp) performance values over GML data

Data Size (KB)	DOM (dom4j)			Pull Parsing
	Document parsing (msec)	Geometry data extraction (msec)	Total time (msec)	Total time (msec)
1	437	32	469	63
10	453	41	494	79
100	547	78	625	172
1,000	625	135	760	234
5,000	1,093	329	1,422	500
10,000	2,094	763	2,857	890
100,000	40,266	∞	∞	7,156
150,000	383,829	∞	∞	70,610
200,000	∞	∞	∞	98,343
500,000	∞	∞	∞	152,172
1,000,000	∞	∞	∞	165,610




Figure 5: Performance comparison of two XML data processors, pull parsing and Document Object Model by using dom4j.

1.3.2. Federator (AWMS)-oriented Approaches

The proposed federated GIS system inherently enables load balancing and parallel processing and this helps in enhancing the overall system performance. This chapter presents the techniques and system design to develop high performance federated GIS system.

The proposed system design and applied techniques are based on WMS Aggregator (AWMS). AWMS is actually a WMS [WMS] with some extensions providing enhanced map rendering services by using innovative pre-fetching, parallel processing and caching techniques. AWMS aggregates, composes and orchestrates WMS and WFS services and, express the layer level compositions in its capabilities file by federating other services' capabilities metadata, and present it to the users through "GetCapabilities" Web Service interface.

System design changes depending on the characteristics of the data to be handled for performance considerations. For the un-frequently changing data we propose pre-fetching (Chapter 1.3.2.1) technique. For the frequently changing data we propose hybrid approach composed of caching and parallel processing techniques applied together (Chapter 1.3.2.2).

In summary, pre-fetching is purely for overcoming the natural bandwidth problem, caching helps system prevent redoing the jobs of querying and rendering before, and parallel processing helps workload sharing and parallel job run. Depending on the data characteristics, AWMS uses only one or the combination of these techniques. These techniques will be explained in the following sections.

1.3.2.1. Pre-Fetching

In the proposed integration framework we deal with archived data provided by WFS in GML format. Archived data does not change often. So, it is not reasonable transferring and rendering the same data again and again for every request coming from the different or even the same users. In order to solve this problem we use pre-fetching. Pre-fetching is used to overcome the performance degradation of transferring large sized data transfer from WFS to WMS. It can even be used in data transfer between WFSs in case of cascaded WFS services.

Pre-fetching is briefly defined as getting the data before it is needed. We accomplish the pre-fetching by the data transfer technique explained in Section 1.3.1.1. The general architecture for the pre-fetching is shown at Figure 6. A performance result of the pre-fetching and comparisons to the on-demand fetching techniques are displayed in Figure 7 and Figure 8 respectively. Since it is independent of real-time application and run in asynchronous manner, it does not degrade the

proposed framework's overall performance. It's running times defined by the periodicity parameter of the Pre-fetching module (PM) (see Figure 6).

The OGC's standard WMS and WFS specifications are based on HTTP Get/Post methods, but this type of services have several limitations such as the amount of data that can be transported, the rate of the data transportation, and the difficulty of orchestrating multiple services for more complex tasks. Web Services help us overcome some of these problems by providing standard interfaces to the tools or applications we develop.

As in proposed data exchange framework defined in Section 1.3.1.1 pre-fetching module make the requests with standard SOAP messages but for retrieving the results a NaradaBrokering subscriber class is used. Through the "*getFeature*" interface of WFS Web Services, pre-fetching module gets the topic name (publish-subscribe for a specific data), IP and port on which WFS streams the requested data. Second request is done by NaradaBrokering Subscriber using the returned parameters. GML data is provided by streaming WFS (implemented by G. Aydin) [Vretanos]. It uses standard SOAP messages for receiving queries from the clients; however, the query results are published (streamed) to a NaradaBrokering topic as they become available. In order to do that, we define the "task" and "timer". Task defines pre-fetching job, and timer defines the running periodicity of the task. Different data might have different periodicities set. Pre-fetching is done over the critical data. Critical data is the GML data affects the performance because of their sizes.

There will be two separate locations for the pre-fetched data. One is temporary which will be active during pre-fetching the data. Another is stable which will be used for serving the clients' requests. When the data transfer is done to the temporary location, all the data at that location will be moved to stable location. Reading and writing the data files at the stable locations will be

synchronized to keep the data files consistent. This cycle is repeated at some time intervals pre-defined by periodicity parameter of Pre-fetching Module (PM).

In order for the pre-fetching algorithm to work properly, pre-fetching module fetches the data as a whole; no constraint should be defined in the query. On the other hand, requests from clients contain some query constraints. These queries and their constraints are handled at the Aggregator WMS (see APPENDIX xx) side. Queries are processed by using parser techniques and XPATH queries over the pre-fetched data.

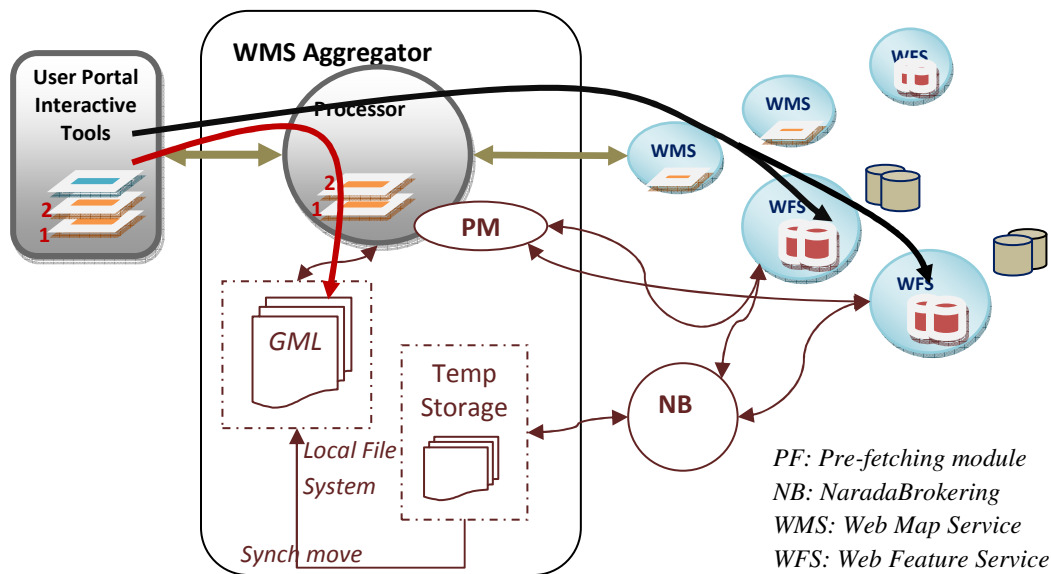


Figure 6: Pre-fetching architecture embedded to the federated GIS system

1.3.2.1.1. Fetching module (PM)

Pre-fetching module composed of two components. One is “timer” defining the periodicity that PF will be running, and other is “task” defining what to do. The periodicity should not be less than the time to transfer one set of critical data. Assigning a periodicity at PM is the most critical

task. This is defined under the considerations of data characteristics and developer's experience with the domain specific application.

Since we develop the system in JAVA we use Timer and TaskTimer JAVA class libraries to implement the routinely running pre-fetching module.

Here is the "task" defined in a pseudo code:

```
public void pseudo_TASK() {  
    Vector CDMdataList = new Vector();  
    CDMdataList = getListPerformanceCritical_GMLDataNames();  
    String tempDatastore = applpath + "/prefetchedData";  
    String stableDatastore = applpath + "/prefetchedDataUsed";  
  
    //Fetching all the data in CDM format (GML) - with NB  
    fd.FetchDataWithStreaming( NBip,NBport,NBtopic,  
                               wfs_address,tempDatastore,CDMdataList );  
  
    //After pre-fetching is done move the data to stable storage  
    fd.moveData(tempDatastore, stableDatastore);  
}
```

We also define timer determining the periodicity of this task to run. Below sample code sets the periodicity of "task" defined above to three days. It means PF will be running once every three days.

```
Timer timer = new Timer();  
  
timer.schedule(task, 0, 40000);
```

Timer class schedules the specified task for repeated fixed-delay execution, beginning after the specified delay. Subsequent executions take place at approximately regular intervals separated by the specified period.

There are two concerns in developing an efficient pre-fetching architecture. First one is limited storage capacity for a node. The size of the pre-fetched data is constrained by local node's hard-disk capacity. Second one is regarding the pre-fetched data characteristics. Some archived data is updated so often that they look like real-time data. In that case, pre-fetching becomes un-feasible and cannot be utilized. For this type of data (archived but updated frequently), we use parallel processing approach applied together with the caching (see Chapter 1.3.2.2).

1.3.2.1.2. Performance Evaluation

We test the proposed pre-fetching technique over the proposed federated GIS system by using real Pattern Informatics (PI) geo-science applications (see Figure 6). PI is an earthquake forecasting application and uses archived earthquake seismic records stored at WFS as feature collections encoded in GML (XML encoded structured data model for geo-data). As explained before, data is kept at WFS and rendered at WMS to create a comprehensible data in map images.

We basically test the system as illustrated in Figure 6. Red-curve (short) illustrates the round-trip path for the pre-fetch data and black-curve (long) illustrates the round-trip path for the on-demand fetched data. For the simplicity we will be using only one critical data to apply pre-fetching.

In summary, we give the performance results for the proposed pre-fetching approach and compare it with the ordinary on-demand fetching approach (see Figure 8 and Table 5). In case of

on-demand fetching approach, one end is database and other end is user browser (see the black (dark)-curve in Figure 6). Performance results show the response times.

Table 4: Performance results for the pre-fetching technique

GML Data Size in MB	Response time pre-fetched data usage seconds
1	3
20	7
40	10
60	13
80	17
100	19

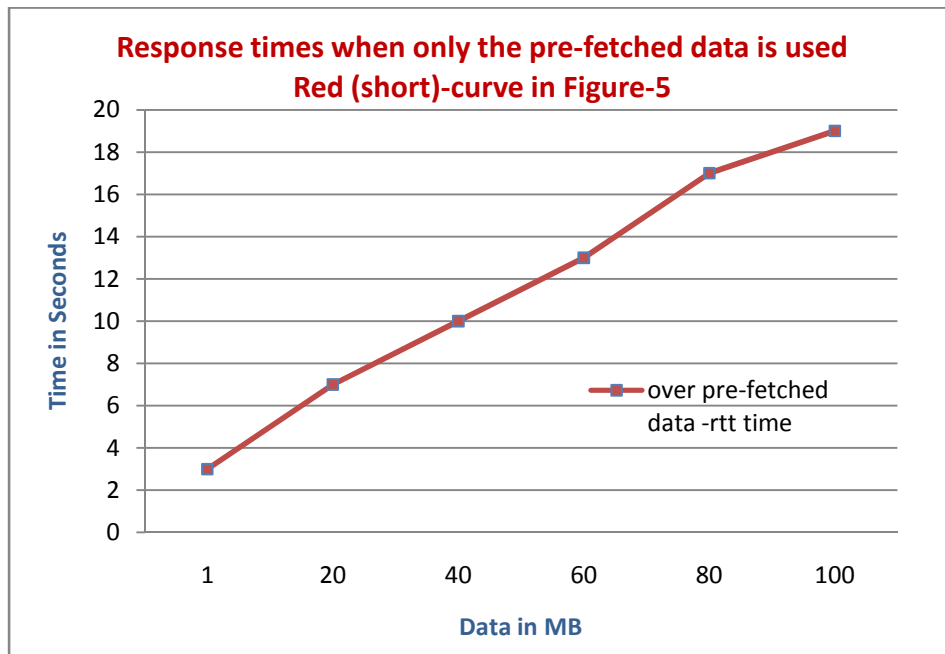


Figure 7: Performance of the pre-fetching technique

Table 5: Comparison of the pre-fetching and ordinary (on-demand fetching) techniques

GML Data Size in MB	with on-demand minutes	with pre-fetching minutes	Log of on-demand minutes
1	1.97	0.05	0.29
20	19.25	0.12	1.28
40	38.17	0.17	1.58
60	58.20	0.22	1.76
80	76.75	0.28	1.89
100	95.80	0.32	1.98




Figure 8: Performance comparison of the map rendering in the proposed GIS system with pre-fetching and ordinary ways.

Figure 8 is actually comparison of *Figure 2* (showing on-demand ordinary system's performance) and *Figure 7* (pre-fetching system's performance). Note that both approaches use the streaming data transfer and advanced pull-parsing techniques. The difference comes only from using the pre-fetching technique.

As it is expected the pre-fetching increased the performance and responsiveness of the system for accessing, querying and rendering archived data. Compared to on-demand fetching (ordinary), pre-fetching removes the times spent on conversion (from database to GML at WFS side) and transferring GML data. In case of cascaded data (going through multiple chained services to access the original data source), performance gains even becomes much larger.

1.3.2.2. Locality-based Query Decomposition and Parallel Processing

This chapter presents another federator (AWMS) oriented high performance design for accessing and rendering of XML encoded large size data (GML).

Parallel processing is implemented based on the main query partitioning. Each partition is assigned to separate thread of work. The number of partitions and their sizes are defined by using locality principles. Locality information is obtained from the cached data kept for the same session and user. In order to achieve this, browser-based sessions are mapped to service-based sessions by updating the headers of the SOAP messages. All the services in the system are Web Services and they communicate through SOAP messages.

AWMS apply the parallel processing after cached data extraction and rectangulation. Since all the data in the system is geo-referenced and queried in ranges defined by bounding boxes (defining coordinates of rectangles in the form of (minx, miny, maxx, maxy)), we do range query partitioning to implement parallel processing. Partitioning the queries for parallel processing can be considered in two ways. One is layer-based (horizontal) another is partitioning the bbox criteria of the queries. These issues are taken care of at the AWMS. For the vertical partitioning,

AWMS defines the layer compositions for hierarchical data (map images) in its capability file. Vertical partitioning is handled in the capabilities file by defining different service access points for different layers. The techniques explained here is regarding the horizontal partitioning. The horizontal partitioning is explained as range query partitioning to assign the main job to worker nodes justly.

Caching + parallel processing algorithm has three parts closely related to each other:

1. *Caching*
2. *Cached-data extraction and Rectangulation*
3. *Parallel-processing*

In order to make these concepts more clear let's illustrate these in a picture (See Figure 9): Figure shows a map image composed of two layers. One is NASA satellite base map layer, and other is a layer showing earthquake seismic records (in blue dots). (a) shows partially overlapping of cached data and main request bboxes. (b) Shows cached data extraction and rectangulation for the remaining part in the main request. (c) Shows partitioning of the rectangles from (a) based on the locality information obtained from the cached data. All the rectangulated regions from (c) will be assigned to a thread to created map images as final responses.

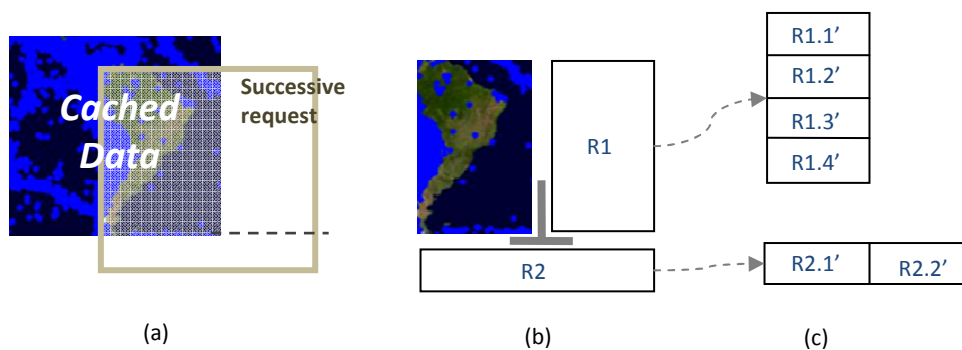


Figure 9: (a) Cached data extraction, (b) rectangulation, and (c) partitioning for parallel processing.

Research questions are summarized as below:

- i. What to cache, how to cache and how to use
- ii. How a server knows what requests come from what user to utilize cache
- iii. How to map browser-based sessions across the servers.
- iv. How to determine the number of the partitions required for the parallel processing
- v. How to partition the rectangles (ex:R1 and R2)
- vi. How to make locality-based query decomposition.
- vii. How to create sub-queries for the partitioned regions (rectangles)
- viii. How to assemble the results to sub-queries to create final response

These research questions are answered in the following two sub-sections. Chapter 1.3.2.2.1 explains the caching, cached data extraction and *rectangulation* processes as answers to i, ii, iii. Chapter 1.3.2.2.2 explains the parallel processing based on the range query partitioning and the locality principles as answers to iv, v, vi, vii, viii.

1.3.2.2.1. Caching, Cached-data Extraction and Rectangulation

We first explain the caching policy and techniques. Second, we introduce a novel approach for mapping browser-based session context to server-wide sessions in order to apply locality principles. Third, we explain techniques for cached data extraction and rectangulation.

1. Caching:

We apply the caching for only the critical data pre-defined in AWMS. Critical data is vector data encoded in GML common data model.

Caching will be utilized just by the successive requests for the same user and session. For each separate session even from the same user there will be separate cached data kept at AWMS. To

do that, we introduce browser-like sessions for each user across the servers. In other words, cached data is kept till the next request comes from the same session.

2. *Mapping Browser-based Session across the Servers.*

Since the proposed federated GIS is interacted through interactive decision making tools over the integrated data views, the session tracking and transfer issues have to be addresses to handle locality based query decomposition and parallel processing across the servers. Java Server Pages (JSP) defines session ID whenever a user opens a page to interact with the federated GIS system. A session is normally stored in a cookie which is available to all windows in the browser. The system access this ID by *session.getId()*. This returns a string unique user ID (uuid) which can be used for application specific purposes.

Whenever federated GIS client interacts with the system through AWMS, it sets its browser's session ID to the header of its SOAP messages sent to the Web Service. All the requests coming from the same browser has same session ID. Session IDs are created when the browser is opened and kept same until it is closed. Each browser has a separate and unique session ID. By setting this session ID to the header of SOAP messages AWMS can distinguish what client (browser) makes the requests and check its cached data and session information stored before.

Here is the pseudo code briefly explaining the steps:

```
WMSServicesSoapBindingStub binding;  
  
binding = (WMSServicesSoapBindingStub)  
        new WMSLocator().getWMSServices(new URL( service_address));  
  
String sessionID = session.getId();  
  
String channel_name = "WMS_getMap_Request";
```

```
//Add SessionID to the SOAP message's header  
binding.setHeader(service_address, channel_name, sessionID);  
  
//See Appendix xx for the sample GetMap request  
  
Object value = binding.getComprehensibleData(getMap_request);
```

Whenever a user access the system through the same browser its session number will be the same and AWMS keeps its local data and actions in the system differentiated based on its unique session ID.

3. Cached Data Extraction and Rectangulation:

According to OGC standards in GIS domain, queries are created with location parameter and location is defined in bounding box (bbox) formats. bbox is a formula defining the region as a rectangle through coordinates of bottom left corner and top right corner. Example: Q(minx, miny, maxx, maxy).

The proposed GIS services are OGC compatible and implemented in Web Service principles. They accept the requests in predefined XML-structured queries such as “getMap” for WMS and “getFeature” for WFS (see Chapter xx). Queries to WMS and WFS are actually window shape range queries. Range queries are formulated in bbox. After extraction of cached data falling in the main query range, the remaining part needs to be converted to the rectangular shapes in order to create valid sub-queries in the ranges defined by the bboxes. This is why we make rectangulation after cached data extraction from queried-region.

Cached-data extraction methodology is removing the regions in the main query which overlap with the cached data and then, creating the rectangular sub-regions from the remaining main query in the form of bboxes (see Figure 9 -b).

Since we use session-based caching (each session has separate cached data) based on the browser sessions, the position of the bboxes of the cached data and the following request determines our strategy for the cached data extraction and rectangulation. The bbox ranges of the cached data and main query can be positioned to each other in four possible ways.

Main query bbox is described as (minx, miny, maxx, maxy)

Cached data bbox is described as (minx^c, miny^c, maxx^c, maxy^c)

(1) The main query covers cached data (zoom-out action),




Figure 10: zoom-out case of rectangulation

Meeting condition:

if(minx^c>minx and maxx^c<maxx
and miny^c>miny and maxy^c<maxy)

Rectangles:

R1: minx, miny, maxx, maxy R3: minx^c, miny^c, maxx^c, maxy^c
R2: maxx, miny, maxx, maxy R4: minx^c, miny, maxx^c, miny^c

(2) The main query falls in cached data (zoom-in action)




Figure 11: zoom-in case of rectangulation

Meeting condition:

if(minx>minx^c and maxx<maxx^c
and miny>miny^c and maxy<maxy^c)

This case enables the fastest response. There is no need for query partitioning and data transfer from WFSs. It just uses cached GML to create map image based on the bboxes values of main query. A lot of performance gains.

(3) The main query partially overlaps with cached data (move action).

This case is explained in *Figure 9*, a and b.

Here is the formula of the rectangles:

R1: $\text{minx}, \text{miny}, \text{maxx}, \text{miny}^c$

R2: $\text{maxx}^c, \text{miny}^c, \text{maxx}, \text{maxy}$

In this case, there are four different sub-cases depending on the movement directions. These are (1) south-east, (2) south-west, (3) north-east, and (4) north-west. The *Figure 9* illustrates the south-east case, and the rectangles above belong to that case. The rectangles for the other cases are also created similarly.

The rectangles obtained in this section go through the partitioning process explained in the following chapter.

1.3.2.2.2. Parallel-Processing with caching

The proposed parallel processing is based on range-query (defined in bbox) decomposition and throughout the document we call it partitioning. The partitioning is done with the locality principles to share the workload to the threads to reduce the response times to a reasonable level.

Partitioning will be done over the rectangular regions obtained at the end of rectangulation process. Issues to make parallel-processing with the locality principles are summarized as below:

- i. How to determine the number of partitions and their sizes in bbox
- ii. How to partition the rectangles
- iii. How to create sub-queries corresponding to the partitions
- iv. How to assemble the results to sub-queries for the main query

These issues are illustrated in Figure 12 below. In this specific example, main query includes three separate layers, and one of them is created with the critical data encoded in common data model, GML. The regions 1 and 2 in the main query are determined by the cached-data extraction and rectangulation processes explained in chapter 1.3.2.2.1. Grey region in the main query overlaps with the cached data. There is no need for data transfer for this region. This is obtained from the cache. For the other parts not overlapping with the cache (region 1 and 2), the system makes parallel processing for data access, query and plotting after creating partitions.

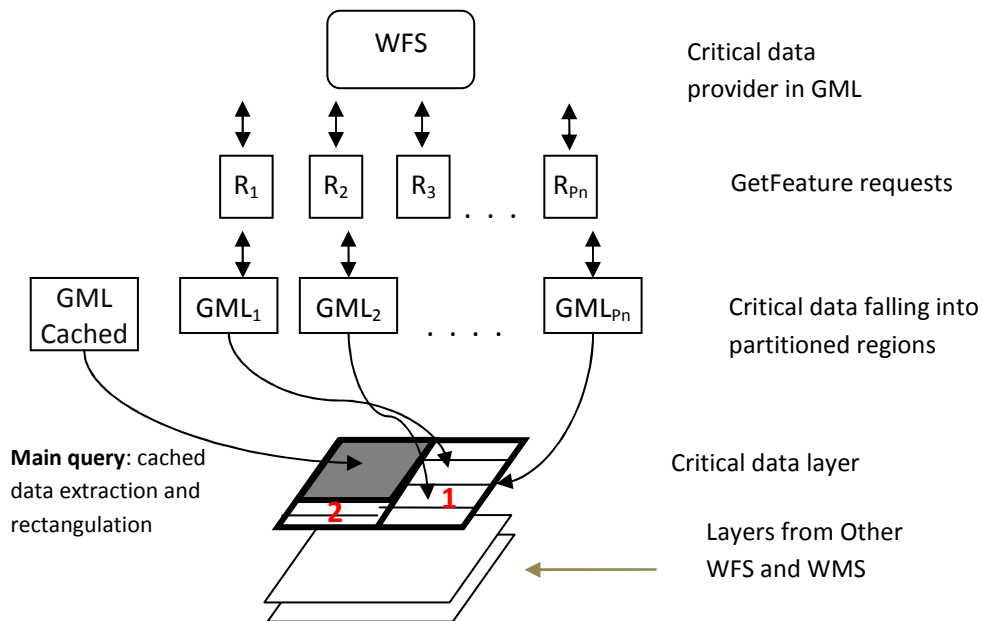


Figure 12: Parallel processing and caching architecture in brief.

For the simplicity, we assume there are one critical data on which parallel processing and caching will be applied on, and one rectangle (obtained through the rectangulation process) to be partitioned in the system.

Terminologies to be used:

- CD_size_kb : Cached-data size measured in KB. Cached-data is kept in local file system as GML file. This is obtained from the system.

- $CD_size_br^2$: Cached data bbox area.

$$CD_size_br^2 = (maxx^c - minx^c)*(maxy^c - miny^c)$$

- $R_size_br^2$: Rectangle's bbox area (ex: Region-1 in critical data layer in Figure 12)

$$R_size_br^2 = (maxx - minx)*(maxy - miny)$$

- Thr : Threshold value. Allowed max data size falling in a partitioned region, and measured in KB.

Threshold value changes from data to data. It is predefined based on the experience of the developer. For example for the earthquake seismic data its value is 1000KB. This value defines a data size for which one thread (ordinary system) can response in a reasonable time period.

- Pn : the number of partition

Let's explain the research issues listed above.

i. Determining the number of partitions:

In order to define the number of partitions we use locality principles. Locality principle in this context is explained as following. If a region has a high volume of data, then the regions in close neighborhood also expected to have high volume of data. Example is the human population data. The urban areas have higher human population than the rural areas. The oceans (2/3 of the world) have no populations etc.

We partition the rectangles into equal regions because we don't know the size of the data falling in that region before getting it. In order to define the size (in bbox) we use cached data sizes expressed in bbox and KB. We assume (by using locality) cached data density is similar to the

main request density, and by using the threshold value and un-cached main request part we calculate P_n as below:

$$\text{Density of cached-data: } dcd = \frac{CD_size_kb}{CD_size_br^2}$$

$$\text{Allowable largest area to assign: } lat = \frac{\text{threshold data size}}{\text{density of cached data}} = \frac{thr}{dcd}$$

$$P_n = \frac{R_size_br2}{\text{largest area to assign}} = \frac{R_size_br2}{lat}$$

If P_n is less than 1 then, don't make partition. In contrast, if it is bigger than 1, then partition into P_n regions. ii explains how to partition a rectangle into P_n number of regions.

ii. *Query decomposition of the rectangulated regions with P_n :*

We already know the bbox of the rectangles obtained through rectangulation process explained at Chapter 1.3.2.2.1. We also calculated the number of regions (in item i) into which each one of these rectangles is going to be partitioned as the set of bbox values.

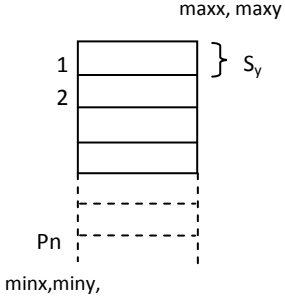
Here, we explain how to partition a given rectangle into P_n number of bboxes. There are two alternative techniques here, one is partitioning the rectangle vertically and the other is partitioning it horizontally.

In case of vertical partitioning the step value is calculated as below, and partitioning is done along the Y-coordinate. (See Figure 13)

$$s_y = \frac{(maxy - miny)}{P_n}$$

In case of horizontal partitioning the step value is calculated as below, and partitioning is done along the X-coordinate,

$$S_x = \frac{(\max x - \min x)}{P_n}$$

 <p>Figure 13: Partitioning a rectangle along the coordinate-y</p>	<p>Calculating the bboxes of the partitioned regions:</p> <p>Vertical:</p> <pre>for (i=0; i<Pn*sy; i=i+sy;) print (minx, miny - i, maxx, maxy-(i+sy));</pre> <p>Horizontal:</p> <pre>for (i=0; i<Pn*sx; i=i+sx;) print (minx-i, miny, maxx-(i+sx), maxy);</pre>
--	---

After having the rectangles partitioned, the partitions go through the query creation process explained as following.

iii. Creating the queries for the sub-regions obtained at ii.

Throughout the rectangulation and partitioning, the only changing attribute of the main query is the bbox coordinate value. These are calculated in the previous step.

Based on the set of bbox values obtained at the end of partitioning process (ii) we need to create sub queries. Each partition is differentiated by only their bbox value, and they go through the query creation process. AWMS creates *getFeature* requests corresponding to these rectangles based on their bounding boxes. Other parameters and attributes required for creating *getFeature* request are obtained from the main query. All the parameters, attributes and their values (except for bbox values) will be the same for all the *getFeature* requests created for the partitions.

An example case of decomposing a rectangle obtained by rectangulation process and creating parallel queries is illustrated at Figure 14. In this example, rectangle is partitioned into 5 regions vertically.

$$P_n = 5 \quad \text{and} \quad s_y = \frac{(\text{maxy}-\text{miny})}{P_n} = \frac{(45-40)}{5} = 1$$

You can also see a sample getFeature created for bbox value “-110, 35, -100, 40” request at Figure 15.

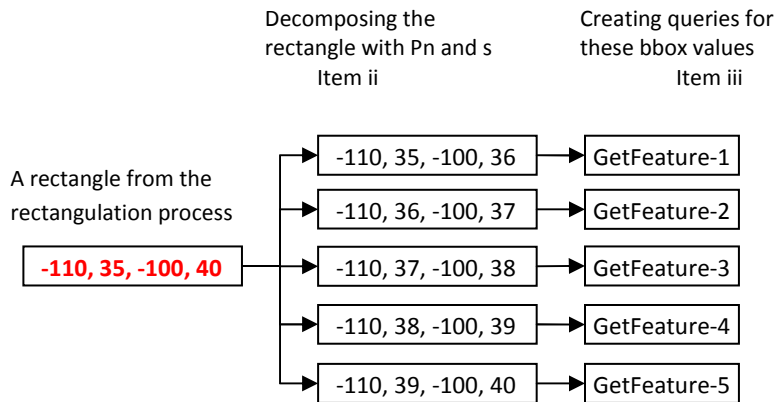


Figure 14: Example scenario of the partitioning a region into 5 sub-regions through the bbox value of a rectangle.

iv. Assembling the results from the sub-queries

Each query created at step iii is sent to WFS in a separate thread, and the returned results are stored locally at AWMS. After getting the data AWMS starts rendering and plotting the critical data over the other layers by parsing and extracting the geometry elements in returned GML. Data transfer, parsing and rendering issues are explained in the previous chapters.

In addition, all the GML data corresponding to main query for the specific bbox is kept as cached data in order to serve the following request coming from the same session.

```

<?xml version="1.0" encoding="iso-8859-1"?>
<wfs:GetFeature      outputFormat="GML2"      xmlns:gml="http://www.opengis.net/gml"
xmlns:wfs="http://www.opengis.net/wfs" xmlns:ogc="http://www.opengis.net/ogc">
  <wfs:Query typeName="global_hotspots">
    <wfs:PropertyName>YEAR</wfs:PropertyName>
    <wfs:PropertyName>MONTH</wfs:PropertyName>
    <wfs:PropertyName>DAY</wfs:PropertyName>
    <wfs:PropertyName>LATITUDE</wfs:PropertyName>
    <wfs:PropertyName>LONGITUDE</wfs:PropertyName>
    <wfs:PropertyName>MAGNITUDE</wfs:PropertyName>
    <ogc:Filter>
      <ogc:BBOX>
        <ogc:PropertyName>coordinates</ogc:PropertyName>
        <gml:Box>
          <gml:coordinates>-110,35-100,40</gml:coordinates>
        </gml:Box>
      </ogc:BBOX>
    </ogc:Filter>
  </wfs:Query>
  <wfs:Query typeName="global_hotspots">
    <ogc:Filter>
      <ogc:PropertyIsBetween>
        <ogc:Literal>MAGNITUDE</ogc:Literal>
        <ogc:LowerBoundary>
          <ogc:Literal>7</ogc:Literal>
        </ogc:LowerBoundary>
        <ogc:UpperBoundary>
          <ogc:Literal>10</ogc:Literal>
        </ogc:UpperBoundary>
      </ogc:PropertyIsBetween>
    </ogc:Filter>
  </wfs:Query>
</wfs:GetFeature>

```

Figure 15: Sample GetFeature request for the partitioned region of bbox (-110, 35 -100, 40).

1.3.2.2.3. Performance Evaluation

Performance will be evaluated in three possible generalized situations categorized based on the cached data utilization. These are (as explained in different context in Chapter 1.3.2.2.1):

- a. No usage of cached data
- b. Complete usage of cached data. Cached data is enough to respond to main request
- c. Partial usage of cached data.

a. **No cached data utilization:**

This case happens when the successive query is randomly created and does not overlap with the cached dat. In this case there is no cached data extraction and sub-rectangle created by rectangulation processes. There is only one rectangle which is the main query. In order to make performance evaluations, we test the system with different levels of partitions such as 2, 4, 6, 8, 10, 15 and 20, and assign them to separate individual threads for the parallel processing. The measured response times are in seconds. Queries are decomposed based on their bbox values. Query bbox values and corresponding data sizes are displayed in Table 6.

Table 6: The performance results for the selected bbox values and corresponding data size for the various levels of parallel processing.

Sample Bounding boxes's of requests				GML Data	Parallel processing - partition numbers -seconds						
minx	miny	maxx	maxy	Size -MB	2thrd	4	6	8	10	15	20
-119.97	36.95	-118.55	38.27	1.081	62	53	43	38	31	30	21
-120.67	36.3	-117.85	38.92	1.711	88	60	45	49	51	41	32
-122.08	34.99	-116.44	40.23	4.999	170	136	108	81	64	329	86
-124.85	32.26	-113.56	42.75	11.110	434	313	295	190	162	144	154
-130.5	27.01	-107.92	48	14.799	558	452	435	315	235	202	178
-141.78	16.52	-96.62	58.49	17.511	668	515	500	512	485	281	242
-164.36	-4.46	-74.04	79.47	34.128	1035	723	656	611	576	540	402
-180	-46.42	-28.88	90	62.696	2608	1208	954	857	922	966	707
-180	-90	46.67	90	83.167	3194	1969	1188	1223	1293	1476	1158

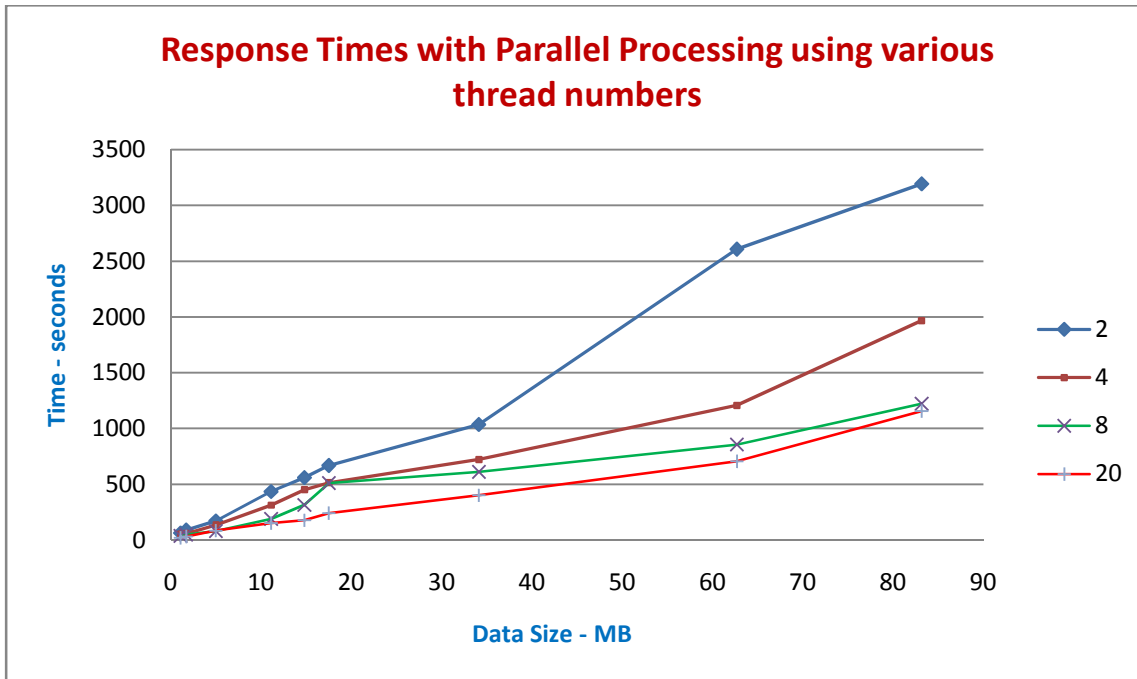


Figure 16: Response times in seconds for different levels of threaded run (2, 4, 8, 20), and data sizes.

Table 7: Comparing response times of 20-threaded parallel processing and single threaded ordinary system.

GML Data Size -MB	Response times -seconds	
	Parallel Proces 20 Threads	ordinary One-threaded
1.081	21	66.670
1.711	32	102.100
4.999	86	292.950
11.110	154	636.280
14.799	178	847.500
17.511	242	1001.920
34.128	402	1957.620
62.696	707	3590.720
83.167	1158	4819.860

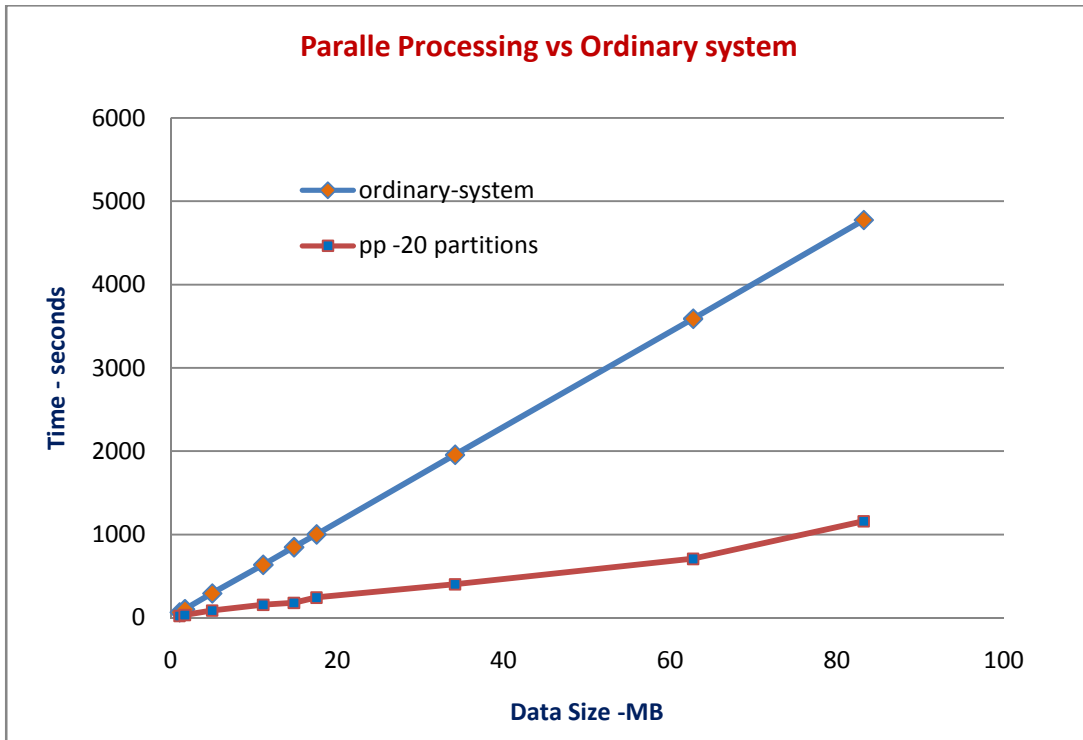


Figure 17: Performance comparison, parallel processing vs ordinary system.

The test results show that 20-threaded parallel processing for the data handling is almost 4-times faster. However, the performance does not increase in the same ratio at which the thread number increases. That is because of the overheads resulted from mainly the query decomposition and assembling the result sets for the main query etc. Moreover, the figure shows that the higher the data size the larger the performance gains.

b. Complete cached data utilization

This case happens when the user query a smaller region falling in the previous map he got on his browser. It is mostly caused by zooming-in action. In this case, the cached data is enough for responding to the main request, and no other cascaded requests are needed. AWMS renders the map just by using the cached data. The only task needed is the cached data extraction and

overlaying (plotting) over the other requested layers. There is no need for rectangulation, the main query decomposition and threaded parallel processing.

Table 8: Response times in case of complete use of cached data. No other requests are done for the response.

Sample Bounding boxes's of requests				Cached GML	Response-time
minx	miny	maxx	maxy	Data Size -MB	seconds
-119.97	36.95	-118.55	38.27	1.081	1.38
-120.67	36.3	-117.85	38.92	1.711	1.45
-122.08	34.99	-116.44	40.23	4.999	1.81
-124.85	32.26	-113.56	42.75	11.110	2.83
-130.5	27.01	-107.92	48	14.799	3.58
-141.78	16.52	-96.62	58.49	17.511	3.92
-164.36	-4.46	-74.04	79.47	34.128	6.62
-180	-46.42	-28.88	90	62.696	11.45
-180	-90	46.67	90	83.167	13.27




Figure 18: Responsiveness of the system in case of complete cached data utilization. Enough cached-data.

Table 9: The response times for complete cached-data and one-threaded ordinary system (Figure 2).

Cached GML Data Size -MB	Complete caching Response time seconds	Complete caching Response time seconds	Ordinary systm Log(Rtt_ordinary) minutes
1.081	1.38	0.02	0.14
1.711	1.45	0.02	0.16
4.999	1.81	0.03	0.26
11.110	2.83	0.05	0.45
14.799	3.58	0.06	0.55
17.511	3.92	0.07	0.59
34.128	6.62	0.11	0.82
62.696	11.45	0.19	1.06
83.167	13.27	0.22	1.12

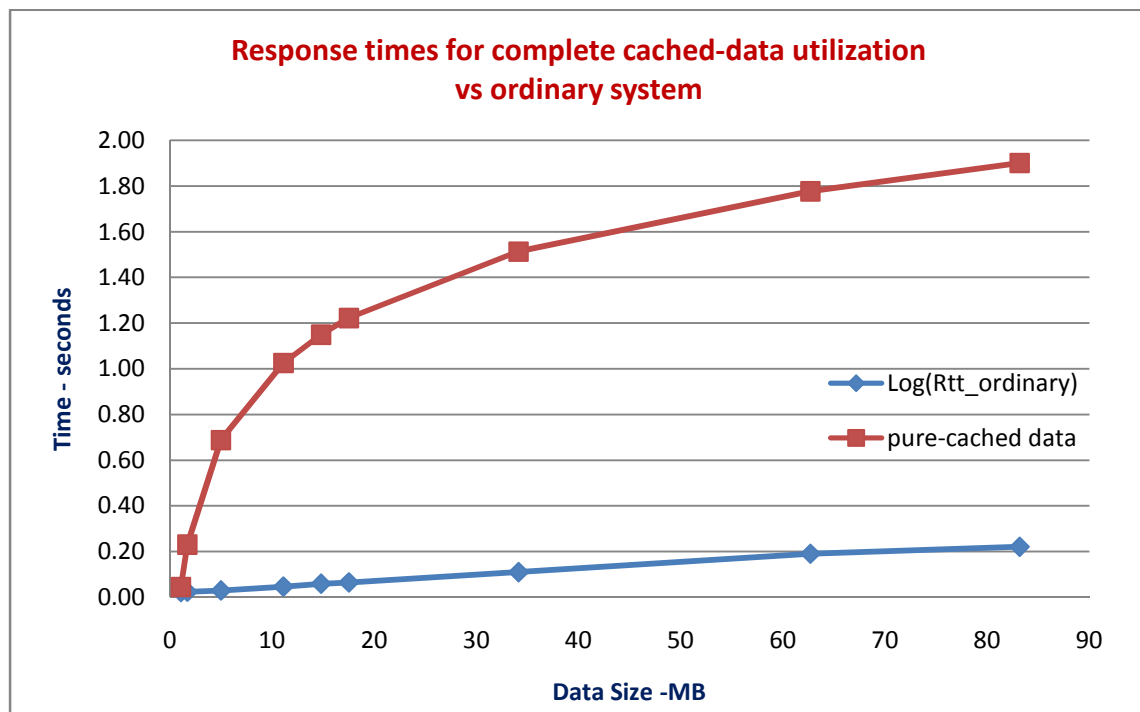


Figure 19: Comparisons of roundtrip times of complete use case of cached data and ordinary systems.

Table 9 and Figure 19 compare the performance results obtained by using complete caching and the ordinary one-threaded data handing. Since most of the time is spent on data transfer and converting any data to common data model GML (at WFS), getting rid of these burdens by using cached data enhances the performance a lot. Note that, in order to show huge performance difference in one graph clearly, the ordinary system performance results are adjusted by taking their logarithmic values.

c. Partial cached data utilization:

This case happens when the user moves (or drag and drop) the map to another region or makes zooming-out. As explained before (*Figure 9* and Chapter 1.3.2.2.1), if only one point of main request falls in bbox boundaries of the cached-data, they are called as partially overlapped. Table 10 lists the sample request bboxes and their corresponding data sizes.

Table 10: Performance results for the sample case scenario in which half of the data is provided by the cached data, and other half is requested by again 20-thread parallel processing.

Sample request distinguished in their bboxes					Parallel Processing	Parallel processing + Caching		
Bounding boxes				Data Size		1/2 cached -seconds		
minx	miny	maxx	maxy	MB	20 thread	time-cd	time-fd	time-sum
-119.97	36.95	-118.55	38.27	1.081	21	0.97	10.50	11.47
-120.67	36.3	-117.85	38.92	1.711	32	1.18	16.00	17.18
-122.08	34.99	-116.44	40.23	4.999	86	1.31	43.00	44.31
-124.85	32.26	-113.56	42.75	11.110	154	1.62	78.00	78.62
-130.5	27.01	-107.92	48	14.799	178	1.91	89.00	90.91
-141.78	16.52	-96.62	58.49	17.511	242	2.82	121.00	123.82
-164.36	-4.46	-74.04	79.47	34.128	402	3.33	201.00	204.33
-180	-46.42	-28.88	90	62.696	707	5.18	353.50	358.68
-180	-90	46.67	90	83.167	1158	7.34	579.00	586.34

Table 10 also presents the performance values for the sample case scenario in which main request bbox values partially overlap with the cached data and half of the data size corresponds to main request is provided by the cached data. For the comparison reasons, un-cached data will be fetched and rendered through 20-thread parallel processing (see Figure 20).




Figure 20: The performance gaining over 20-threaded pp when half of the data is fetched from the cache.

Figure 21 and Table 11 compares the performance results of the hybrid system (caching+parallel processing) with ordinary on-demand system. caching+parallel processing they use the values obtained at the sample scenario shown in Table 10.

As it is shown the in Figure 21, proposed system is almost 10 times faster than the ordinary on-demand one-threaded system. As the data size increases, that ratio increases.

Table 11: Comparison of response times for caching+pp and ordinary non-caching single-threaded system

Data Size MB	Response times -seconds	
	1/2 cached time-sum	ordinary system
1.081	11.47	66.670
1.711	17.18	102.100
4.999	44.31	292.950
11.110	78.62	636.280
14.799	90.91	847.500
17.511	123.82	1001.920
34.128	204.33	1957.620
62.696	358.68	3590.720
83.167	586.34	4819.860




Figure 21: Comparing response times of caching+pp and ordinary system.

REFERENCES

- [Lu 2006] Wei Lu, Kenneth Chiu, and Yinfei Pan, “A Parallel Approach to XML Parsing”.
In *the 7th IEEE/ACM International Conference on Grid Computing, 2006*.
- [Pallickara2003] Pallickara S. and Fox G., “NaradaBrokering: A Distributed Middleware Framework and Architecture for Enabling Durable Peer-to-Peer Grids”
ACM/IFIP/USENIX, Rio Janeiro, Brazil June 2003.
- [Donbox] Don Box, David Ehnebuske, Gobal Kakivaya, Andrew Layman, Dave Winer.,
Simple Object Access Protocol (SOAP) Version 1.1, May 2000.
- [Sosnoski] Sosnoski, D. “XML and Java Technologies”, performance comparisons of the
Java based XML parsers. Available at [http://www-
128.ibm.com/developerworks/xml/library/x-injava/index.html](http://www-128.ibm.com/developerworks/xml/library/x-injava/index.html)
- [Alexander] Aleksander Slominski. XML Pull Parser, visited 04-15-02.
<http://www.extreme.indiana.edu/xgws>.
- [Vretanos] Vretanos, P. (ed.), Web Feature Service Implementation Specification (WFS)
1.0.0, OGC Document #02-058, September 2003
- [GML] Cox, S., Daisey, P., Lake, R., Portele, C., and Whiteside, A. (eds) (2003), OpenGIS
Geography Markup Language (GML) Implementation Specification. OpenGIS project
document reference number OGC 02-023r4, Version 3.0.
- [WMS] de La Beaujardiere, J., Web Map Service, OGC project document reference number
OGC 04-024. 2004.
- [WFS] Vretanos, P. (2002) Web Feature Service Implementation Specification, OpenGIS project
document: OGC 02-058, version 1.0.0. Volume,

[Booth]Booth, D., Haas, H., McCabe, F., Newcomer, E., Champion, M., Ferris, C., and Orchard, D. "Web Service Architecture." W3C Working Group Note, 11 February 2004.

Available from <http://www.w3c.org/TR/ws-arch>.

[Tran] Tran, P., Greenfield, P., and Gorton, I., *Behavior and Performance of Message-Oriented Middleware Systems*. . Proceedings of the 22nd international Conference on Distributed Computing Systems, ICDCSW. 2002.

[uddi] Bellwood, T., Clement, L., and von Riegen, C., UDDI Version 3.0.1: UDDI Spec Technical Committee Specification <http://uddi.org/pubs/uddi-v3.0.1-20031014.htm>. 2003.