TELEIOS

FP7-257662

Deliverable

7.4

An evaluation of the real-time fire monitoring application and the TELEIOS infrastructure

Charalabos (Haris) Kontoes, Ioannis Papoutsis,
Dimitrios Michail, Themos Herekakis

and

Consortium members

31/08/2013

Status: Final

Scheduled Delivery Date: 31 August 2013

Executive Summary

The main purpose of this document is to present the main results from the evaluation of the real-time fire monitoring system that was developed in the framework TELEIOS, with the collaborative effort of WP7, WP1, WP4, WP5 and WP2 activities. This system comprises of a fully automatic fire monitoring processing chain that combines EO image acquisitions, with auxiliary geo-information and human evidence in order to draw reliable decisions and generate highly accurate fire products.

Database and semantic technologies that were developed in WP2, 4 and 5, and were integrated in WP1 which were combined to offer a state-of-the-art solution for effectively monitoring active fire-fronts of forest wildfires are also evaluated.

Additionally, the development of a new SciQL-based service, presented in D7.5 and D5.3, to map burnt areas after forest wildfires is also evaluated, and the potential scalability of the service to cope with an increased number of Landsat imagery is discussed.

Document Information

Contract Number	FP7-257662	**Acronym**	TELEIOS
Full title	Virtual Observatory Infrastructure for Earth Observation Data		
Project URL	http://www.earthobservatory.eu/		
EU Project officer	Francesco Barbato		

Deliverable	**Number**	D7.4	**Name**	An evaluation of the real-time fire monitoring application and the TELEIOS infrastructure			
Task	**Number**	T7.4	**Name**	Evaluating the real-time fire monitoring application and the TELEIOS infrastructure			
Work package	**Number**	WP7					
Date of delivery	**Contractual**		M36 (August 2013)	**Actual**	31 August 2013		
Status	draft ☐ final ☑						
Nature	Prototype ☐ Report ☑						
Distribution Type	Public ☑ Restricted ☐ Consortium ☐						
Authoring Partner	National Observatory of Athens (NOA)						
QA Partner	German Aerospace Center (DLR)						
Contact Person	Dr Charalabos (Haris) Kontoes						
	Email	kontoes@space.noa.gr		**Phone**	+30-2108109186	**Fax**	+30-2106138343

Project Information

This document is part of a research project funded by the IST Programme of the Commission of the European Communities as project number FP7-257662. The Beneficiaries in this project are:

Partner	Acronym	Contact
National and Kapodistrian University of Athens Department of Informatics and Telecommunications (Coordinator)	NKUA	Prof. Manolis Koubarakis National and Kapodistrian University of Athens Dept. of Informatics and Telecommunications Panepistimiopolis, Ilissia, GR-15784 Athens, Greece Email: koubarak@di.uoa.gr Tel: +30 210 7275213, Fax: +30 210 7275214
Fraunhofer Institute for Computer Graphics Research	Fraunhofer	Dr. Eva Klien Fraunhofer Institute for Computer Graphics Research Fraunhofer Strasse 5, D-64283 Darmstadt, Germany Email: eva.klien@igd.fraunhofer.de Tel: +49 6151 155 412, Fax: +49 6151 155 444
German Aerospace Center The Remote Sensing Technology Institute Photogrammetry and Image Analysis Department Image Analysis Team	DLR	Prof. Mihai Datcu German Aerospace Center The Remote Sensing Technology Institute Oberpfaffenhofen, D-82234 Wessling, Germany Email: mihai.datcu@dlr.de Tel: +49 8153 28 1388, Fax: +49 8153 28 1444
Stichting Centrum voor Wiskunde en Informatica Database Architecture Group	CWI	Prof. Martin Kersten Stichting Centrum voor Wiskunde en Informatica P.O. Box 94097, NL-1090 GB Amsterdam, Netherlands Email: martin.kersten@cwi.nl Tel: +31 20 5924066, Fax: +31 20 5924199
National Observatory of Athens Institute for Space Applications and Remote Sensing	NOA	Dr. Charalambos Kontoes National Observatory of Athens Institute for Space Applications and Remote Sensing Vas. Pavlou & I. Metaxa, Penteli, GR 152 36 Athens, Greece Email: kontoes@space.noa.gr Tel: +30 210 8109186, Fax: +30 210 6138343
Advanced Computer Systems A.C.S S.p.A	ACS	Mr. Ugo Di Giammatteo Advanced Computer Systems A.C.S S.p.A Via Della Bufalotta 378, RM-00139 Rome, Italy Email: udig@acsys.it Tel: +39 06 87090944, Fax: +39 06 87201502

Table of Contents

1 Introduction

1.1 Purpose of this Document

The main purpose of this document is to present the results from the thorough evaluation of two services developed in the framework of TELEIOS, namely the real-time fire monitoring and the Burnt Scar Mapping (BSM) services, which are elaborated in D7.3 “Implementing the real-time fire monitoring application – Phase II” [**1**].

The developed system encompasses the semantic and database technologies of WP4 ‘Scalable storage and query processing for EO image metadata’ and WP5 ‘Ad-hoc and continuous/stream queries for EO images’ respectively, as described in [**2**]. The integration of this state-of-the-art expertise, customized to serve the use case needs, was developed in WP1 ‘The TELEIOS infrastructure’ [**3**].

The two services based on Earth Observation data and remote sensing processing techniques are evaluated herein in terms of thematic accuracy and processing times to meet specific requirements. Additionally, the TELEIOS infrastructure is evaluated in terms of a series of criteria pertinent to usability of the system and other functional and non-functional requirements that were defined in the beginning of the project.

It should be mentioned that considering the reviewers’ recommendations for (a) demonstrating scalability with a EO-based application and (b) using input data sets in non-proprietary formats such as the open GeoTIFF or NetCDF formats, and the stimulations from the continuous interaction with the user community, a new fire mapping service has been developed, though not initially included in the TELEIOS Description of Work [**2**]. This service, named Burn Scar Mapping, aims at automatically detecting burnt areas and producing damage assessment maps for (i) Emergency Support GMES activities, (ii) seasonal estimation of the damaged areas, and (iii) generate diachronic damage assessment maps which can be exploited for further scientific analysis by foresters.

1.2 Structure of this document

The rest of the document is organized as follows.

Section 2 briefly goes through the main characteristics of the services that were developed within TELEIOS, the real-time fire monitoring and the burnt scar mapping.

Section 3 is the core part of the deliverable which focuses on the thorough evaluation of the services and TELEIOS components in terms of the thematic accuracy of the products that can be generated, the processing times for the timely response of the system and the conformity to several acceptance criteria.

Section 4 presents the main points that were extracted from the system’s stakeholders during the three TELEIOS dedicated user workshops.

Section 5 discusses the lessons learnt from the application of semantic and array database technologies for ‘real-life’ processing scenarios, along with envisaged future plans in the same direction.

Sections 6, 7, 8 and 9 are dedicated to the list of abbreviations, the references, the list of figures and the list of tables respectively of the present document.

2 Products and services developed in TELEIOS

The TELEIOS advancements to today's state of the art in EO data processing are shown graphically with yellow color in Figure 1 and can be summarized as follows:

1. Traditional raw data processing is augmented by *content extraction* methods that deal with the specificities of satellite images and derive image descriptors (e.g., spectral characteristics of the image). Knowledge discovery techniques combine image descriptors, image metadata and auxiliary data (e.g., GIS data) to determine concepts from a domain ontology (e.g., forest, lake, fire, burned area, etc.) that characterize the content of an image [**4**].
2. Hierarchies of domain concepts are formalized using OWL ontologies and are used to annotate standard products. Annotations are expressed in RDF and are made available as linked data so that they can be easily combined with other publicly available linked data sources (e.g., GeoNames, LinkedGeoData, DBpedia) to allow for the expression of rich user queries.
3. Web interfaces to EO data centers and specialized applications (e.g., hotspot detection) can now be improved significantly by exploiting the semantically-enriched standard products and linked data sources made available by TELEIOS. For example, an advanced EOWEB-NG-like interface to EO data archives can be developed on top of a system like Strabon[1] to enable end-users to pose very expressive queries. Rapid mapping applications can also take advantage of rich semantic annotations and open linked data to produce useful maps even in cases where this is difficult with current technology. Open geospatial data are especially important here. There are cases of rapid mapping where emergency response can initially be based on possibly imperfect, open data (e.g., from OpenStreetMap) until more precise, detailed data becomes available.

In all of the above processing stages, from raw data to application development, TELEIOS utilizes database technologies as Figure 1 illustrates.

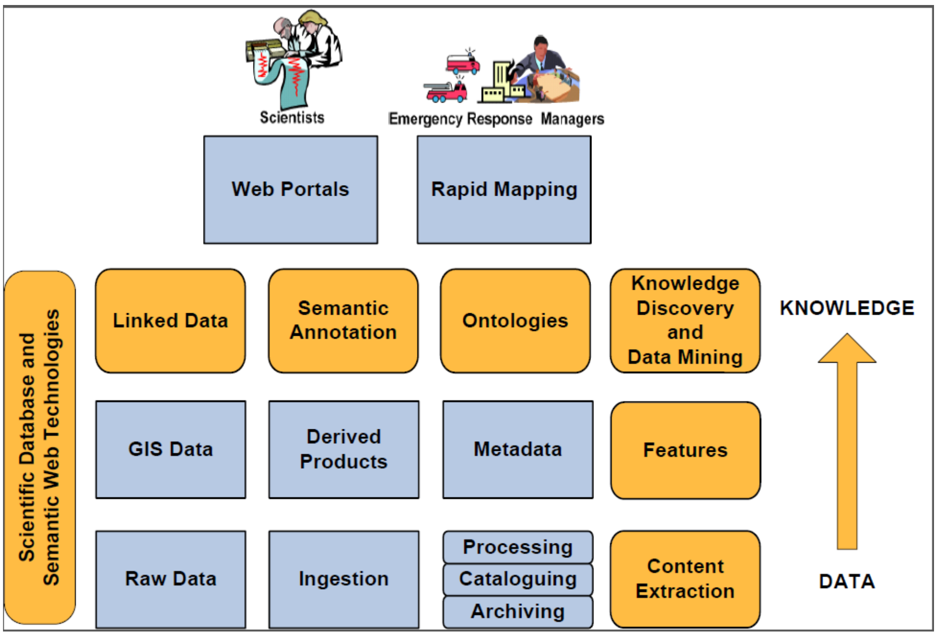

Figure 1. Concept view of the TELEIOS Earth Observatory

[1] http://www.strabon.di.uoa.gr/

The design and implementation of the WP7 use case was based on the functional and non-functional requirements [**5**] collected during the 1st TELEIOS user workshop [**6**], and the feedback gathered in the framework of the 2nd TELEIOS user community workshop [**7**]. The diverse nature of the users that attended the workshops allowed drawing guidelines from both the EO and IT communities in the development of the fire monitoring platform. In [**1**], a thorough description of the service chains that were developed in the framework of TELEIOS can be found, however, for reasons of completeness, a brief description of the services is included hereinafter.

2.1 Real-time fire monitoring

The fire monitoring service is active in NOA before TELEIOS, and since early 2007. It was based on the integration of in-house tools to detect fire hotspots in the entire Greek territory, using a modified version [**8**] of the algorithm in [**9**] and customized detection thresholds to fit the land cover specificities of Greece. The service takes as input MSG-SEVIRI data from Meteosat-8, -9, and -10 satellite platforms and generates hotspots within which it is highly probable that there is a fire.

The fire pixels derived by the above processing chain have dimensions equal to the sensor's spatial resolution, in this case nearly 4x4 km. Thus, MSG/SEVIRI is a low resolution observational system, compared to other very high resolution sensors with similar fire detection capabilities (e.g., WorldView-2 at 0.5 m, Quickbird at 2.4 m, IKONOS at 4 m or Formosat-2 at 8 m), high resolution sensors (e.g., Spot-5 at 10 m and Landsat-5 TM at 30 m), or medium resolution sensors (e.g., MODIS Terra and Aqua with 2 bands at 250 m, 5 bands at 500 m and 29 bands at 1 km). However, the unique advantage of MSG/SEVIRI is its geostationary orbit, which allows for a very high observational frequency (5-15 minutes) over the same area of interest. Other satellite platforms with better spatial resolution are forced to undertake orbits that are closer to the earth, which considerably reduces their revisit time. For example, Aqua MODIS, with its near-polar orbit, passes over Greece twice a day (at 00:30 and 11:30) and the same applies for Terra MODIS (at 9:30 and 20:30). Another important advantage of the MSG/SEVIRI sensor is that its sensitivity is not at all affected by its low spatial resolution, i.e., it is not necessary for an entire 4x4 km pixel to be ‘on fire’ in order to detect the corresponding hotspot. A small pixel portion, exhibiting increased temperature due to a wildfire, will suffice.

However, the service presented a few thematic drawbacks that are discussed in [**1**], and were related to cases of hotspots occurring in the sea, cases of hotspots located outside forested areas, and spatial and temporal inconsistencies in the final product. Moreover, the dispersion of the various processes of the fire monitoring service in many machines and pieces of software made it difficult for NOA to keep all functionalities synchronized. These issues were dealt by using TELEIOS technologies. Figure 2 depicts the new fire monitoring application of NOA, developed in TELEIOS.

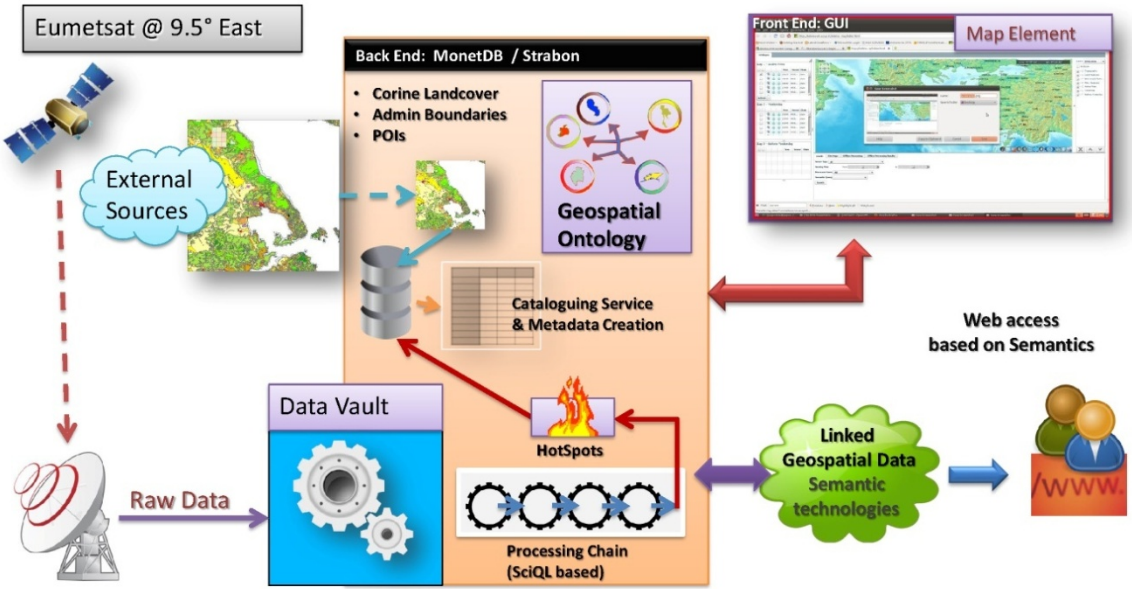

Figure 2. The improved fire monitoring service

The system consists of the following parts:

1. The data vault which is responsible for the ingestion policy and enables the efficient access to large archives of image data and metadata in a fully transparent way, without worrying for their format, size and location.
2. The back-end of the system. The back-end relies on MonetDB for two tasks: (i) the implementation of the hotspot detection processing chain (using the SciQL front-end) [**10**] and (ii) the evaluation of semantic queries for improving the accuracy of the product shapefiles and generating thematic maps (using a stSPARQL front-end, i.e., Strabon) [**11**].
3. A geospatial ontology which links the generated hotspot products with stationary GIS data (Corine Land Cover, Coastline, Greek Administrative Geography), and with linked geospatial data available on the web (LinkedGeoData, GeoNames). This ontology is expressed in OWL.
4. The front-end interface, for controlling the back-end functionality with user-friendly tools, and disseminating the products to the end-user community. A visual query builder is currently being developed as well to allow NOA personnel to express complex stSPARQL queries [**12**].

The processing chain comprises of the following submodules: (a) ingestion, (b) cropping, (c) georeference, (d) classification, and (e) output generation. All submodules are implemented inside the MonetDB DBMS using SciQL. Figure 3 shows an overview of the fire monitoring processing modules. More details can be found in [**1**].

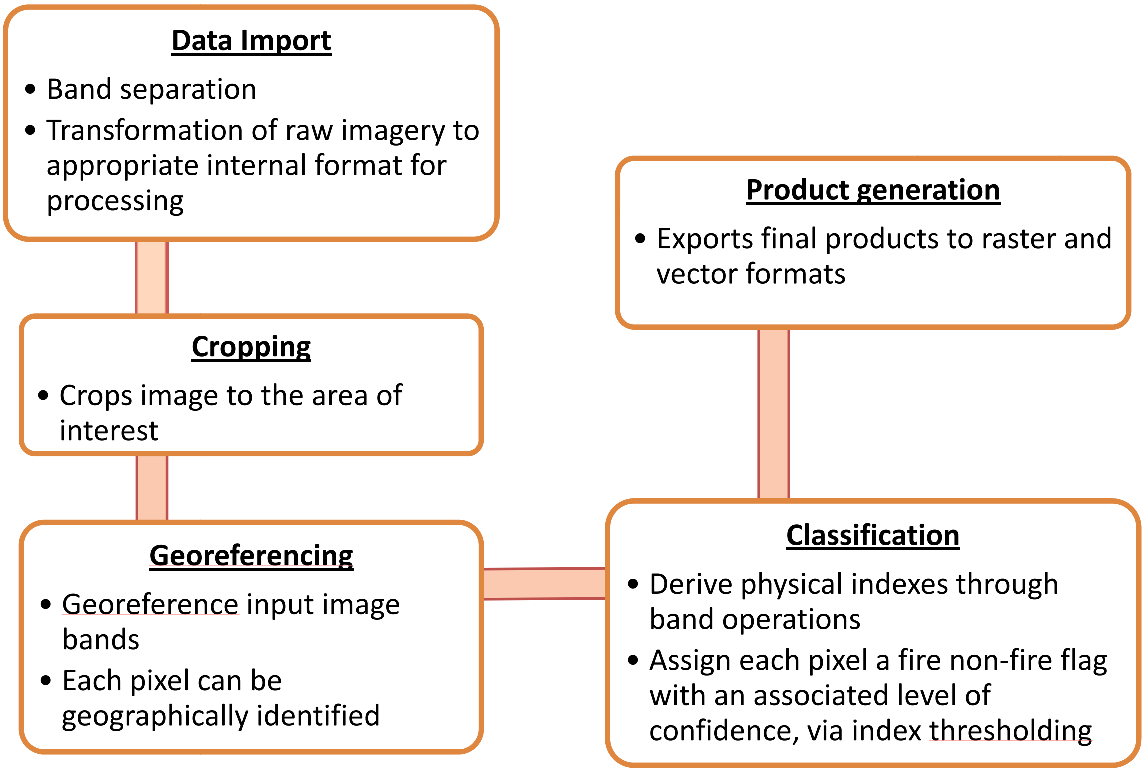

Figure 3. Overview of the fire monitoring processing modules

Most of the fire detection algorithms using MSG/SEVIRI data are based on variations of EUMETSAT's (the international organization managing the Meteosat series of geostationary meteorological satellites) proposed classification methodology for identifying hotspots [**9**]. NOA, among others [**13**], has adopted a similar approach [**8**] for detecting hotspots. However, different studies have been conducted [**14**] and validated [**15**] that suggest another direction. The main concern with respect to EUMETSAT's recommendation has to do with the variance of the solar contribution to earth's apparent temperature, leading to unreliable comparison with predefined thresholds. Therefore, it is suggested that a contextual analysis should be carried out through a spatial matrix of N×N pixels. NOA on the other hand has adopted a dynamic threshold estimation approach, as opposed to fixed thresholding, to cope with the solar variance, leading to increased thematic accuracy [**16**]. It should be noted that [**14**] was applied in the framework of SAFER (EC/GMES) project with noticeable omission errors for Greece, mainly attributed to insufficient customization for the geographic area's special characteristics in terms of vegetation species and underlying land cover.

In the framework of TELEIOS, NOA has been focusing on enhancing the performance of the adopted processing chain in terms of producing results with increased thematic and geometric accuracy, and reliability. Therefore, three main advancements were accomplished: i) applied a new resampling polynomial to achieve better pixel geolocation within the georeferencing module, ii) introduced an extra classification criterion to detect and eliminate false alarms that appear near the edges of specific type of clouds, and iii) adopted an enhanced classification algorithm to avoid omission errors by introducing a second classification approach, where the thresholds are dynamically calculated for each new image acquisition and for every pixel of this image.

In TELEIOS, standard products produced by the SciQL-based processing chain of EO data centers can be combined with auxiliary data to offer to users specific functionalities that go beyond the ones currently available to them. Therefore, this was implemented for the real-time fire monitoring use case, to improve the outputs of the hotspot detection processing chain. An ontology for annotating NOA standard products was designed, and auxiliary geospatial datasets (Coastline of Greece, Greek Administrative Geography[2], CORINE Land Cover[3], LinkedGeoData[4] focusing on publishing OpenStreetMap[5], and GeoNames[6]) are utilized in the fire monitoring application. Then, stSPARQL queries are applied that improve the accuracy of NOA standard products.

For this reason, we are also utilizing the model stRDF, an extension of the W3C standard RDF that allows the representation of geospatial data that changes over time [**11**]. stRDF is accompanied by stSPARQL, an extension of the query language SPARQL 1.1 for querying and updating stRDF data. stRDF and stSPARQL use OGC standards (Well-Known Text and Geography Markup Language) for the representation of geospatial data [**17**].

The database and semantic technologies for the fire monitoring use case, described in [**1**] were integrated to shape the entire service architecture. The back-end and the front-end of the 1st version of the TELEIOS system are extensively described in [**12**], while its evolution to the 2nd and final version is elaborated in [**3**]. A snapshot of the front-end of the TELEIOS system is presented in Figure 4. It incorporates some basic functionalities of a typical GIS environment (zoom, pan, identify tool, etc).

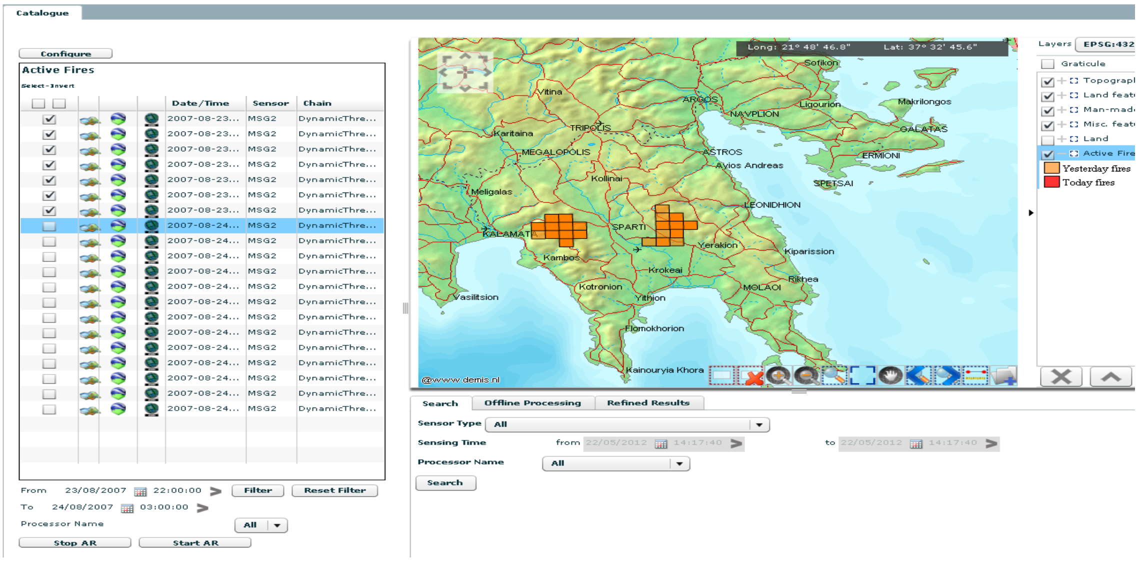

Figure 4. Front-end interface of the 1st version of the TELEIOS system

2.2 Burn Scar Mapping with Landat

An additional application was introduced in TELEIOS, aiming at enhancing NOA's processing capabilities, and highlighting the scalability of the TELEIOS technologies to engage with high data volumes. This application is the BSM, and is dedicated to the

[2] http://geodata.gov.gr/
[3] http://www.eea.europa.eu/publications/COR0-landcover/
[4] http://linkedgeodata.org/
[5] http://www.openstreetmap.org/
[6] http://www.geonames.org/

accurate mapping of burnt areas in Greece after the end of the summer fire seasons, using Landsat 5 TM satellite images.

Since 2007 every September NOA has been producing burnt area maps and delivering them to institutional end-users (GSCP, HMOA/DGF) to include them in their decision making processes [**18**]. Figure 5 depicts the main steps of the BSM production chain.

Figure 5. The BSM service chain

The processing chain is divided into three stages, each one containing a series of modules, namely the pre-processing, the core-processing, and the post-processing phases. In TELEIOS, the Classification module and parts of the Noise removal module were automated using Scientific Python[7] and were then translated into SciQL queries.

A thorough description of the SciQL-based queries can be found in [**1**] and in [**19**]. Synoptically, the chain includes the image loading, the core classification algorithm to burnt and non-burnt areas, the majority filter, and noise removal modules.

3 Evaluation of TELEIOS

In this section the evaluation of the use case application and the TELEIOS system is presented, based on both qualitative criteria and quantitative measures. While Section 3.3 provides a direct link with the acceptance criteria set in [**5**], Section 3.1 attempts to test the enhancement of the thematic accuracy achieved, and Section 3.1.2 highlights the results of the experiments conducted to evaluate the execution processing times using the SciQL version of the hotspot detection chain and the BSM chain, and the performance of the stSPARQL refinement queries.

3.1 Thematic accuracy of fire-monitoring products

3.1.1 Validation of the real-time fire monitoring service

Evaluating the thematic accuracy of the real-time fire monitoring products is not a straightforward task, as it entails the cross-validation of the hotspot products with

[7] http://www.scipy.org/

ground-truth data. The latter is impossible to acquire for the large scale monitoring scenario applicable in this use case. Therefore, we adopted another approach by estimating the relevant thematic accuracy of the MSG/SEVIRI hotspot products with respect to a similar EO application.

Hotspot products are also generated by the MODIS sensor on board of the Terra and Aqua satellite platforms. The spatial resolution is four times better than that of MSG/SEVIRI (1×1 km, as opposed to 4×4 km respectively), but with much less temporal resolution (4 passes over Greece per day, as opposed to the 5 minute acquisition resolution respectively).

The methodology followed for the comparison of the two hotspots products is the following:

1. Selection of an appropriate time span to run the experiments. This was selected to cover three full days in 2007 (24/08, 25/08, 26/08), when Greece was struck by the most severe forest wildfires of the last 20 years.
2. Acquisition of the MODIS hotspots products that served as the reference dataset. The data were collected by FIRMS[8] a NASA portal that integrates remote sensing and GIS technologies to deliver global MODIS fire locations and burned area information.
3. Georeferencing of the reference dataset to the MSG/SEVIRI product projection system.
4. Preparation of the MSG/SEVIRI dataset to be evaluated. Since the temporal resolution of the two datasets is different, we merged 30 minutes of MSG acquisitions (maximum three hotspot products of the MSG2 sensor) around the corresponding MODIS acquisition time.
5. Performing some vector manipulation functions (from point to polygon and vice versa) in order to reliably estimate the overlapping regions of the two hotspot products.

Table 1. Evaluation of the MSG/SEVIRI hotspots derived by the single set of thresholds approach

Timestamp	Time	Total No of FIRMS hotspots	Total No of MSG hotspots	No of FIRMS detected by MSG	Omission error (%)	No of MSG detected by FIRMS	False alarm rate (%)
24/8/2007	0:21	41	25	34	17.07	18	28.00
24/8/2007	9:46	92	116	84	8.70	83	28.45
24/8/2007	11:25	113	120	89	21.24	87	27.50
24/8/2007	20:51	294	155	248	15.65	126	18.71
25/8/2007	1:03	136	59	79	41.91	41	30.51
25/8/2007	8:51	179	234	163	8.94	172	26.50
25/8/2007	12:08	415	467	372	10.36	334	28.48
25/8/2007	19:55	301	185	210	30.23	166	10.27
26/8/2007	0:09	240	83	134	44.17	79	4.82
26/8/2007	9:34	294	179	247	15.99	132	26.26
26/8/2007	11:12	254	285	208	18.11	214	24.91
26/8/2007	20:39	183	8	29	84.15	8	0.00

[8] http://earthdata.nasa.gov/data/near-real-time-data/firms

Total	2542	1916	1897	25.37	1460	23.80

Table 1 and Table 2 summarize the results obtained, when the NOA derived hotspots originate from either the single set of thresholds or the dynamic threshold approach (see Section 2.1) respectively. For each timestamp of the MODIS acquisitions, the corresponding omission error (number of hotspots that were detected by MODIS but not from MSG/SEVIRI) and false alarm rate (number of MSG/SEVIRI hotspots that were not 'seen' by MODIS) are given. The shaded values correspond to these indicators for the full three days of the experiments.

Table 2. Evaluation of the MSG/SEVIRI hotspots derived by the dynamic threshold approach

Date	Time	Total No of FIRMS hotspots	Total No of MSG hotspots	No of FIRMS detected by MSG	Omission error (%)	No of MSG detected by FIRMS	False alarm rate (%)
24/8/2007	0:21	41	60	41	0.00	36	40.00
24/8/2007	9:46	92	116	84	8.70	83	28.45
24/8/2007	11:25	113	120	89	21.24	87	27.50
24/8/2007	20:51	294	287	272	7.48	186	35.19
25/8/2007	1:03	136	216	127	6.62	153	29.17
25/8/2007	8:51	179	234	163	8.94	172	26.50
25/8/2007	12:08	415	467	372	10.36	334	28.48
25/8/2007	19:55	301	442	282	6.31	320	27.60
26/8/2007	0:09	240	257	217	9.58	216	15.95
26/8/2007	9:34	294	179	247	15.99	132	26.26
26/8/2007	11:12	254	287	208	18.11	214	25.44
26/8/2007	20:39	183	92	150	18.03	68	26.09
Total		2542	2757	2252	11.41	2001	27.42

The validation protocol adopted to indentify the omission error is the following: for each timestamp (columns 1 & 2 in Table 1) the points corresponding to the total MODIS hotspot detections (third column in Table 1) was overlaid with the polygons of the total MSG hotspots (fourth column in Table 1). The number of MODIS hotspots falling inside the MSG polygons, with 700 m tolerance (accounting for the 1 km pixel size of MODIS), are registered in the fifth column of Table 1. Then, the omission error (sixth column in Table 1) is calculated by the formula below. A graphical representation of this is seen in Figure 6.

$$\text{Omission error} = \frac{\{(\text{Total No of FIRMS hotspots}) - (\text{No of FIRMS detected by MSG})\}}{(\text{Total No of FIRMS hotspots})} \times 100\%$$

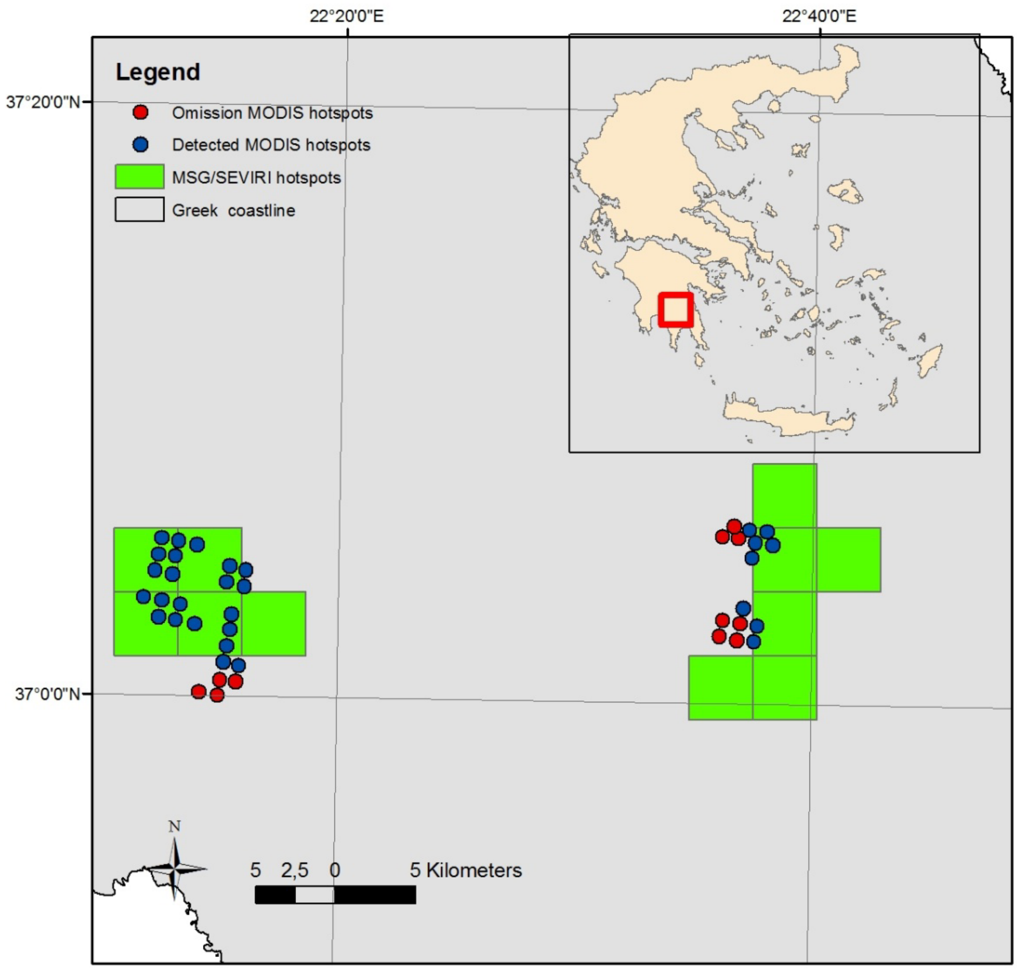

Figure 6. Example of the identification of omission errors for the hotspots of 24/08/2007 (00:21)

A similar procedure was followed for the estimation of the false alarm rate. A buffer of 3.5 km around the MODIS hotspot was applied (accounting for both the MSG and MODIS pixel sizes), and an overlay between this buffered layer and the total MSG hotspots (transformed into a points layers by calculating the centroid of the corresponding polygons) gave the number of MSG hotspots that were detected by MODIS (seventh column in Table 1). Finally the false alarm rate (eighth column in Table 1) was calculated by the formula below. Again, a graphical representation is provided in Figure 7.

$$\textit{False alarm rate} = \frac{\{(\textit{Total No of MSG hotspots})\text{-}(\textit{No of MSG detected by FIRMS})\}}{(\textit{Total No of MSG hotspots})} \times 100\%$$

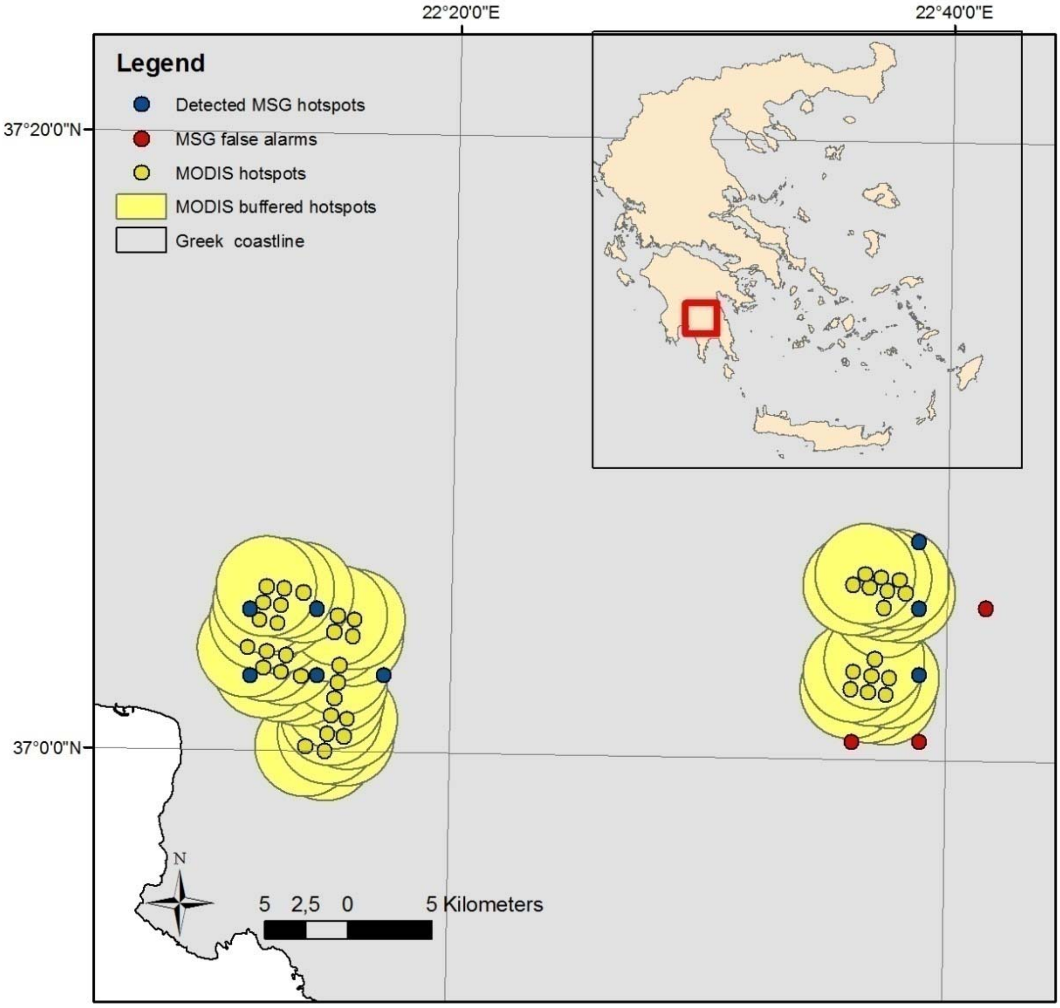

Figure 7. Example of the identification of false alarms for the hotspots of 24/08/2007 (00:21)

The single set of thresholds approach, i.e. thresholds that are constant irrespective of the date & time of the MSG/SEVIRI acquisition and the location of the pixel, achieves an overall relative omission error of 25.37% and false alarm rate of 23.8%. The corresponding values for the dynamic threshold approach are 11.41% and 27.42%, leading to an almost 2.5 times better relative thematic accuracy, at the cost of small increase in the false alarm rate. It should be mentioned that for the selected time span, a total of 2542 hotspots were detected by MODIS, while 1916 & 2757 by MSG/SEVIRI for the static and dynamic threshold approaches respectively.

In general, for both techniques, it was observed that the slightly elevated relative false alarm rate is attributed to the inherent discrepancy in the spatial resolutions of the two sensors and the increased temperature sensitivity of MSG/SEVIRI near intense wildfires. The latter essentially means that the false positives do not occur as isolated hotspots but as pixels classified as fires near neighboring pixels detected as fires by both the MSG/SEVIRI and MODIS sensors. This might be the outcome of hot smoke fumes from nearby fires that contaminate neighboring pixels. The latter are either wrongly detected as fires from MSG/SEVIRI or masked out by the MODIS algorithm. A typical example of such a manifestation is depicted in Figure 8, where the smoke for the fire concentrated on the east (MODIS blue dots) is transferred to the west leading to false alarms by MSG/SEVIRI. In the same figure, an omission error is apparent; there is an undetected MSG/SEVIRI pixel in the center of a burning region.

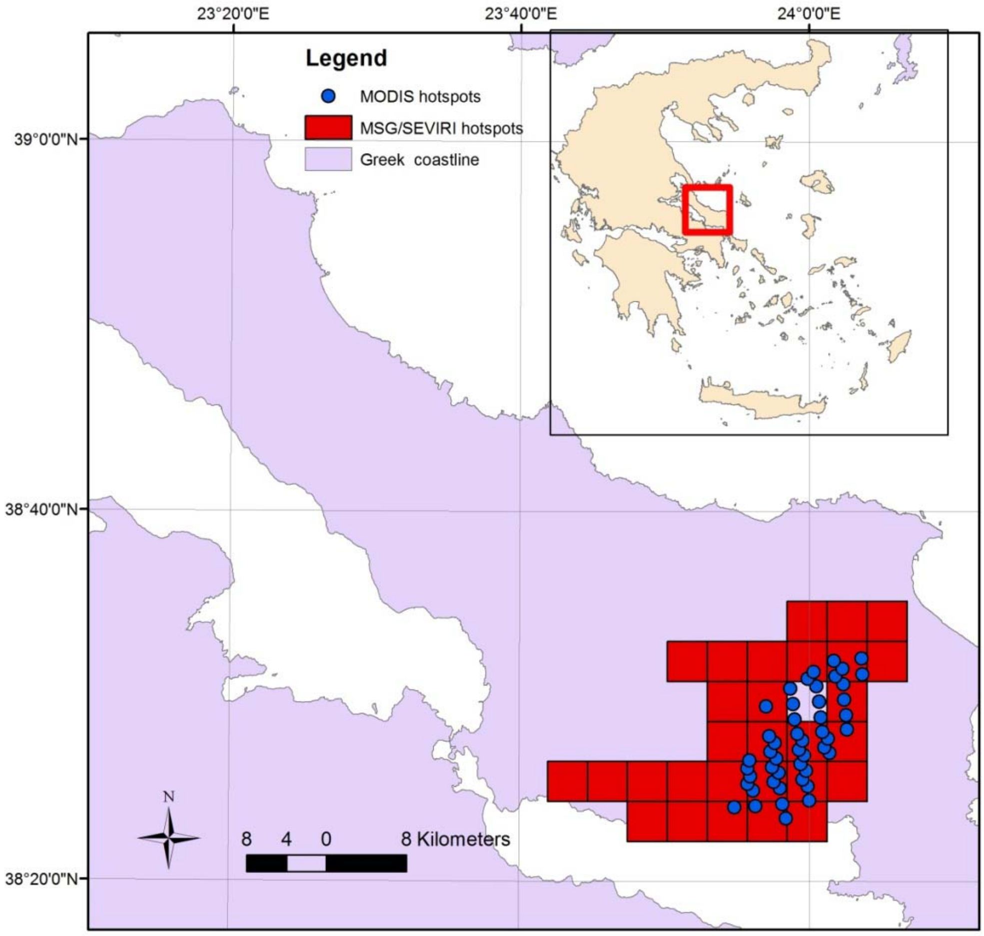

Figure 8. Typical example of false alarms due to smoke and omission error (void rectangle), for the 25/08/2007(12:08) timestamp

Unfortunately MODIS passes over Greece either in the morning or during the night as it can be seen by the available timestamps of Table 2. The enhancement achieved with the dynamic threshold approach can be highlighted though when examining the products during dusk. This is performed in Figure 9, where the improvement in the fire detection capabilities is evident.

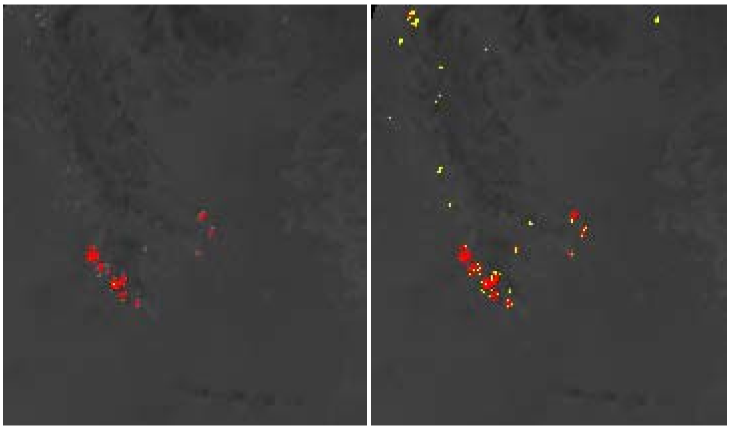

Figure 9. Overview of the forest wildfires for the (left) static and (right) dynamic threshold approaches, for a selected timestamp during dusk (25/08/2007, 18:00). Red pixels correspond to certain fires, while yellow to potential fires with a confidence level of 0.5.

3.1.2 Validation scenarios for BSM

In the framework of the SAFER[9] EC/GMES[10] project, the BSM methodology developed by NOA has been validated by independent evaluators for the needs of the European Commission in a geographic context that excluded Greece. The method's accuracy was assessed over the complex forested landscape of the Corsica Island. In this context, the method was compared in terms of the thematic quality and several other criteria focusing on the assessment of the operational maturity and accuracy standards required by the GMES stakeholders, with other similar methods also used in the framework of GMES for Burn Scar Mapping. The results of this external validation exercise are explicitly stated in [**20**], delivered in the framework of the SAFER project.

The test case was chosen thanks to the ground truth provided by ONF (Office National des Forêts of Corsica). Two fires which occurred in August 2003 in the South of Corsica were considered, in the communes of Tolla and Aullène. ONF provided the ground truth, in terms of fires perimeters. Following this test scenario, the NOA BSM service was qualified –top of its class– as an end-to-end service for fire related Emergency Support activities for integration to operational scenarios all over Europe. The thematic accuracy was originally checked considering the shapefile provided and comparing it to the ONF perimeters. Then all map elements were considered, to confirm the reliability of the information content, the consistency of the information support and the usability of the product.

JRC performed the scientific validation of the NOA BSM product using ground truth fire perimeters and concluded to high producer's and user's accuracies (above 85%) and low commission and omission errors (below 15%). The detailed validation results for the two test areas in Corsica are depicted in Table 3.

[9] http://www.emergencyresponse.eu
[10] http://copernicus.eu/

Table 3. External thematic accuracy assessment of the BSM_NOA chain

Region	**Tolla**	**Aullène**
Commission error	13.10%	5.76%
Omission error	9.32%	12.70%
Producer's accuracy	90.68%	87.30%
User's accuracy	86.90%	94.24%
Fuzzy Kappa	0.843	0.892

Parts of the Python-based processing chain of the BSM methodology were translated to SciQL queries [**1**], [**19**]. The thematic accuracy has not changed in this transition. As an example the burnt areas from the severe 2009 forest wildfire that hit northeast of Athens is presented in Figure 10. No major discrepancies were identified, and only some isolated pixels were not common, mainly attributed to the slightly different implementation of the filtering functionality.

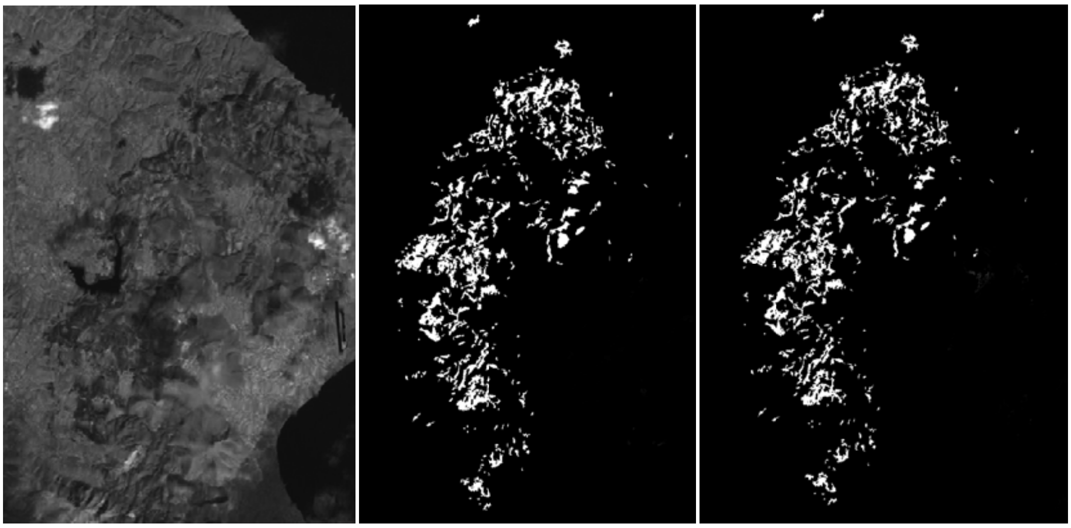

Figure 10. Burn Scar Mapping example with the Python-based (middle) and SciQL-based (right). White areas correspond to pixels classified as burnt areas. The left figure corresponds to Band 4 of Landsat 7 ETM of the area of interest, which is located in northeastern Athens.

3.2 Processing times

In this Section we include the performance evaluation of the current implementation of the two distinct phases of the fire monitoring procedure, a) the fire detection algorithms (Section 3.2.1) and (b) the refinement steps performed using additional information stored in Strabon (Section 3.2.2). Additionally, the results obtained from conducting experiments with the SciQL based burnt scar mapping modules are included (Section 3.2.3).

3.2.1 SciQL based fire-detection algorithm

The original fire detection processing chain has been implemented on top of MonetDB using the array-enabled SciQL language. Two main versions of the processing chain have been already developed, (a) a prototype implementation developed during the first months which used only a subset of the available SciQL language constructs, and (b) a final production ready version which uses all the language features. A major shortcoming of the original prototype is the lack of thread-safety as it is using non-temporary arrays in order to store intermediate results. Thread-safety has been established in the final version since CWI has already implemented the necessary language constructs. In this section we will refer to these two versions as the "prototype" (a) and the "beta" (b) versions. The original NOA implementation of the chain using the C programming language will be referred as the "legacy" implementation.

Table 4. Processing times per image acquisition

	Average Time Per Image (sec)	Max Time Per Image (sec)	Min Time Per Image (sec)
Legacy C Chain	1.481058	1.607081	1.215612
Prototype SciQL Chain	1.795309	2.116046	1.655277
Beta SciQL Chain	2.067308	2.432782	1.902349

The performance of all three versions is compared using all the image acquisitions of a particular day, namely the 22th of August 2010. The dataset for this particular day contains image acquisitions for almost every 5 minute period, i.e., it contains 281 images, most of which contain confirmed fires. All experiments where performed on a Intel(R) Core(TM) i7 CPU 920 processor running at 2.67GHz with 3GB of main memory. The system was running Linux while both MonetDB and the legacy C implementation were compiled with a 32-bit version of GCC 4.4 using the same optimization flags.

The assessment of the three processing chains was performed as if the chains are black boxes, measuring the wall time that each chain requires for each image. Additionally we also recorded the total running time in seconds for the execution of the processing chains. The total running time includes a small additional overhead due to bookkeeping, such as moving around the necessary files, etc. Note that the most expensive part of the processing chain is the decompression of the original satellite image, which is performed in all three cases by the same library. Table 4 presents the minimum, maximum and average time per image that is required for the execution of each chain. Table 5 presents the total time required in order to execute each processing chain for the whole dataset.

Table 5. Total processing time for one day acquisitions

	Total running time for 281 image acquisitions (sec)
Legacy C Chain	468.91
Prototype SciQL Chain	535.51
Beta SciQL Chain	638.68

Figure 11 presents in more detail the running times of all three processing chains executed every 15 minutes in a window of 3 hours between 11am and 3pm. Similarly Figure 12 presents the running times for every 5 minutes in a window of one hour between 12pm and 3pm.

In all 281 timestamps at this particular day of the year, all three processing chains behave similarly. The legacy C chain is the fastest, followed by the prototype implementation and finally followed by the current beta version. This increase in execution time does not pose any significant problems in NOA, as it does not influence in any significant way the amount of time available to the refinement queries. Recall that both the processing chain execution and the refinement queries need to finish in less than 5 minutes. Moreover, NOA expects that the execution time of the latest SciQL chain will improve during the remainder of TELEIOS as the implementation of SciQL matures inside the MonetDB. On the other hand, the use of a high-level scripting language, such as SciQL, for the implementation of the processing chain significantly reduces the development effort. Common small changes, such as changing threshold values, are as easy as changing a few tuples in the DBMS, avoiding lengthy development phases.

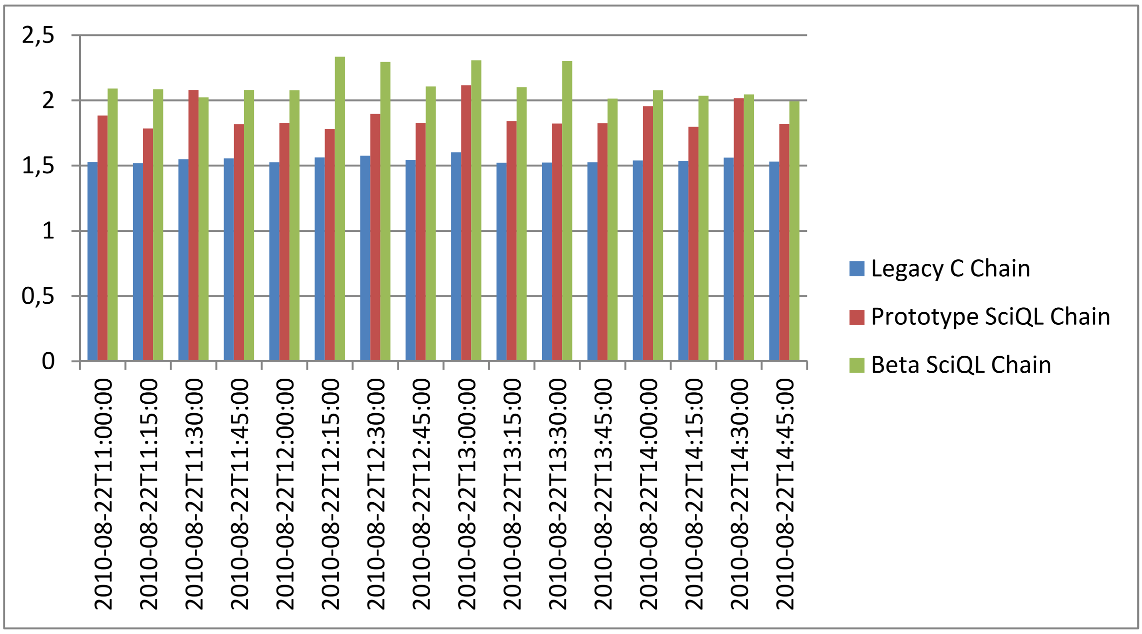

Figure 11. Running time in seconds of the processing chains executed every 15 minutes in a three hour window

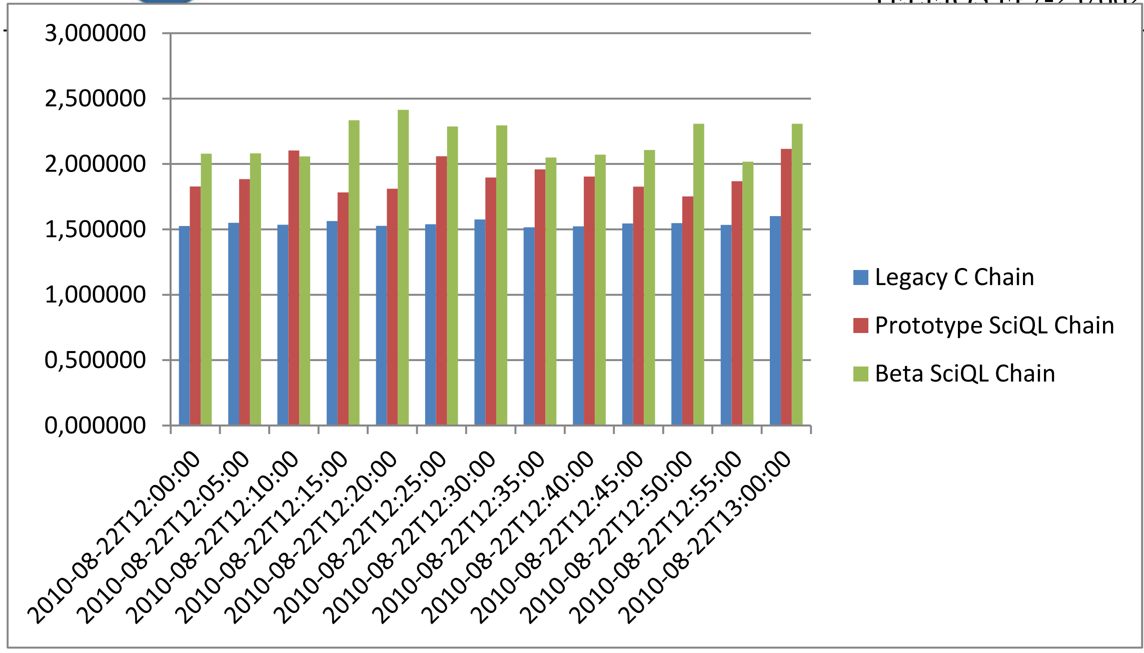

Figure 12. Running time in seconds of the processing chains executed every 5 minutes in a window of one hour

3.2.2 st-SPARQL refinement queries

We have carried out several experiments in different machines for each version of the refinement queries. The datasets we used contained hotspots derived from sensors MSG1 and MSG2 during the fire seasons of the years 2007, 2008, 2010, 2011, and 2012 (up to 19/07/2012), combined with the Greek Administrative Geography dataset and the CLC dataset. The size of the dataset related to hotspot information is around 542,000 triples. The geometry of the Greek coastline was also included so that the respective spatial joins, as we described in the previous sections, could be performed. In this section we present two of our most recent experiments which were executed in one of our machines. The experimental environment and the results are described in the sections that follow.

<u>Experimental Setup</u>

Our experiments were carried out on an Ubuntu 11.04 installation on two Intel Xeon E5620 with 12MB L2 cache running at 2.4 GHz. The system has 48 GB of RAM and 4 disks using RAID configuration as two mirrored sets in a striped set (RAID level 1+0). The metric we used to measure performance is the response time for each query posed for the respective operation by measuring the elapsed time from query submission till a complete iteration over each query's results had been completed.

Experiment Results

For the experiments we presented herein, the following datasets were used:

- Hotspots derived from the sensor MSG1 this year (2012)
- Greek Administrative Geography
- The coastline of Greece
- The Corine Land Cover Ontology

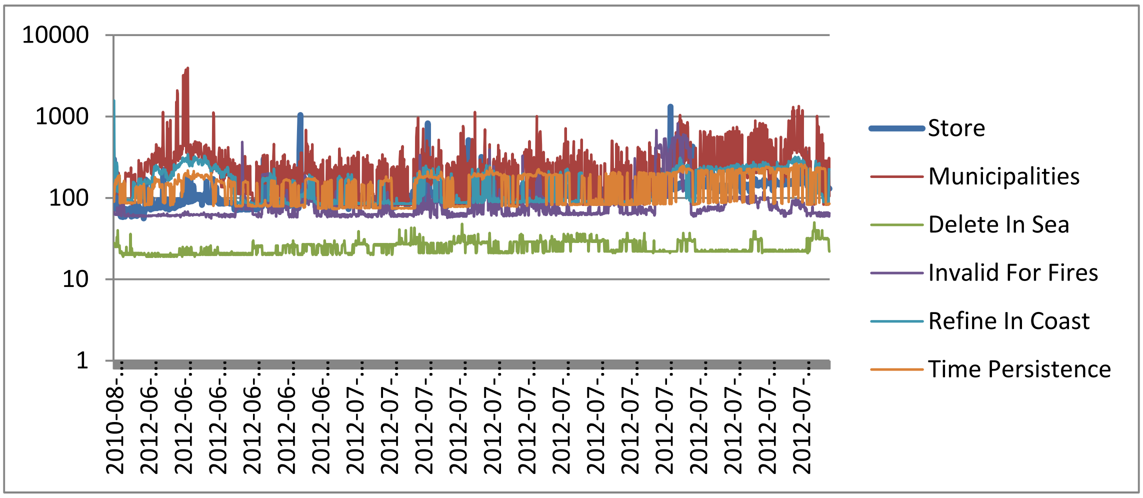

Figure 13. Response time (ms) for each MSG1 acquisition of 2012

The MSG1 sensor detects hotspots every five minutes. The refinement operations described in this chapter were applied to the products of each acquisition and the response time for each operation was measured. The results of this experiment are shown in Figure 13.

In Figure 13, we observe that all operations are executed efficiently, mostly in less than a second, except for the operation of associating each detected hotspot with the municipality it belongs to. This operation is labeled as "Municipalities" in the figure presented above and is shown with orange color, and, although for most cases the query processing time does not exceed two seconds, there are cases where it needs four seconds to be completed. Even in these cases, the performance of the system is satisfactory.

The second experiment we present was executed in the same experimental environment. The datasets used are the following:

- Hotspots derived from the sensor MSG2 during the fire periods of the years 2007, 2008, 2010, and 2011
- Greek Administrative Geography
- The geometry of the coastline of Greece

The MSG2 sensor detects hotspots every fifteen minutes. Similarly, refinement operations are performed to the MSG2 acquisitions and the response time of each respective query is measured, presented in Figure 14.

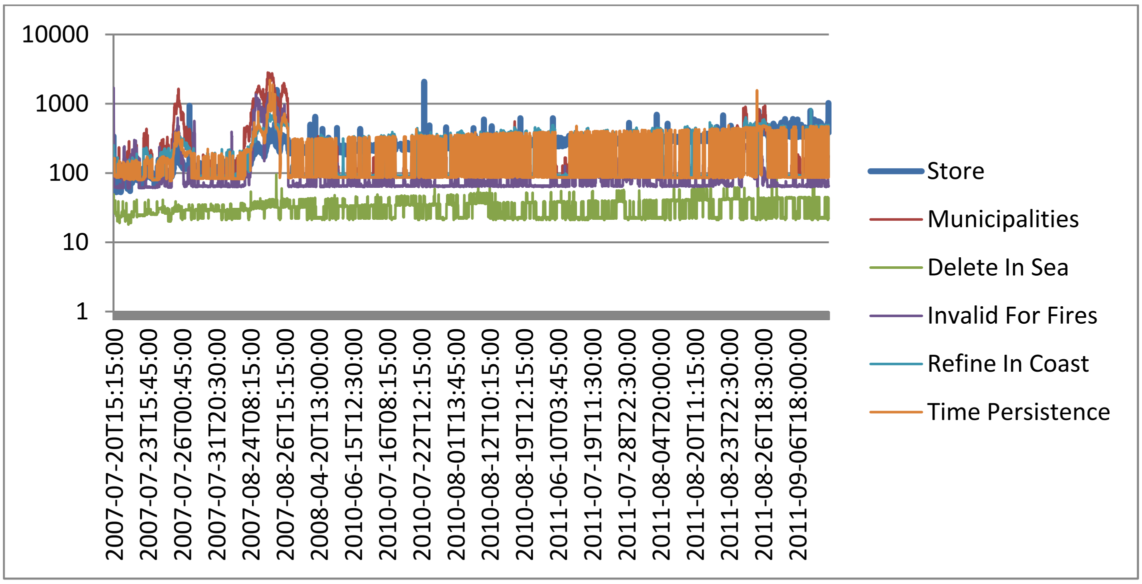

Figure 14. Response time (ms) for each MSG2 acquisition of 2007, 2008, 2010 and 2011

In Figure 14, similarly to Figure 13, we observe some spikes in the response time of Strabon while inserting the Greek Administrative Geography information. An additional observation is that query processing time grows in acquisitions with a larger number of hotspots. After this preliminary evaluation, we observe that the performance of Strabon is satisfactory, given that the sensors MSG1 and MSG2 provide an acquisition every five and fifteen minutes respectively.

3.2.3 SciQL based burnt scar mapping

The burnt scar mapping (BSM) processing chain has been re-implemented in the context of TELEIOS using the SciQL language. In this section we assess this implementation using a large number of input images, comparing it with our original implementation which is done using the Python programming language. The BSM use case is applied on a very large number of images spanning several years of Landsat images.

The performance of the SciQL based BSM and of the Python based implementation is compared using 192 images which cover Greece over a period from 2004 up to 2012. The total size of this archive is 164GB. All experiments were performed on an Intel(R) Core(TM) i7-3770 CPU @ 3.40GHz with 8GB of main memory. The system was running Linux and MonetDB was compiled with a 64-bit version of GCC 4.6 using all available optimization flags. For the experiment we used changeset 578291337244 from the SciQL-2 branch of the public Mercurial MonetDB repository. The MonetDB server was started with the additional options

- gdk_vmtrim=no
- gdk_nr_threads=2

using 2 threads for parallel execution.

The assessment of the processing chains was performed as if the chains are black boxes, measuring the wall time that each chain requires for each image. Table 6 presents the minimum, maximum and average time per image that is required for the execution of each chain.

Table 6 Processing Times per Image.

	Average Time Per Image (sec)	Max Time Per Image (sec)	Min Time Per Image (sec)
SciQL BSM Chain	478	909	57
Python BSM Chain	778	1100	540

The following table presents the total time required to execute the BSM processing chain in the whole archive of images from 2004-2012. We see a noticeable improvement from around 40 hours taken from the original chain to almost 24 hours by the SciQL chain.

Table 7 Total processing time in seconds for processing the whole archive from 2004-2012 which is 192 images. Each image is roughly 770MB.

	Total Time (sec)
SciQL BSM Chain	87933
Python BSM Chain	144583

Greece is covered by images with the following path/row combinations: 180/034-180/035, 181/033-181/036, 182/031-182/036, 183/031-183/035, 184/031-184/034, 185/032-185/033. Figure 15, Figure 16, Figure 17, and Figure 18 present the processing times of the two BSM chains for the above mentioned path/row combinations.

We can easily observe that the SciQL processing chain is always faster than our original Python chain. Note that the original Python chain contains no parallelization while the SciQL chain runs with 2 threads. However, this is the added advantage of using SciQL and a declarative language in order to express the BSM chain. The added benefits of parallelization have been gained without any additional effort from the programmer's side. We expect that as the SciQL implementation in MonetDB matures, additional improvements in the running time will be noticeable.

Figure 15 Processing times for path-rows 180034, 180035, 181033, 181034, 181035 from 2004 up to 2012.

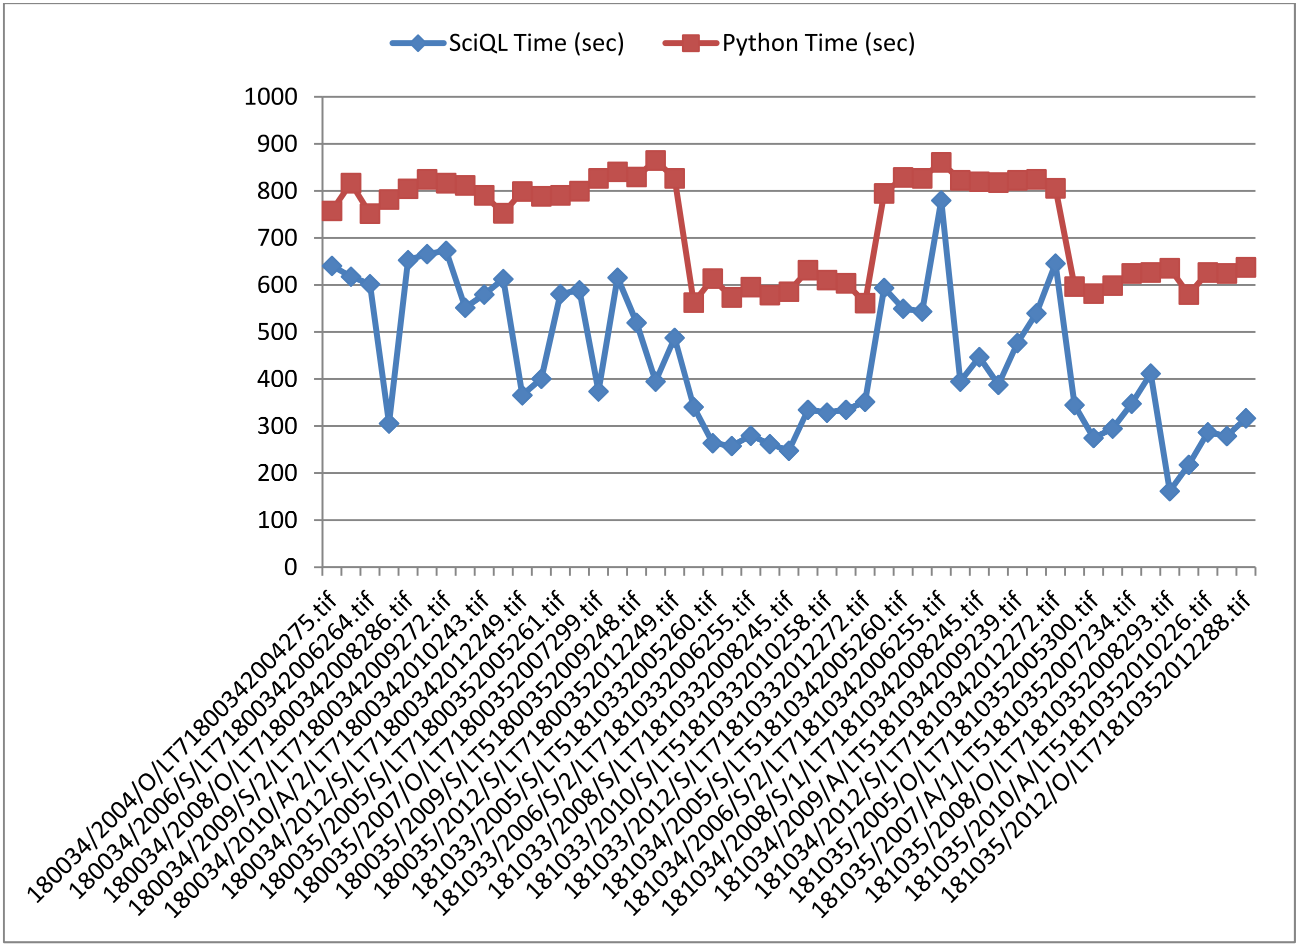

Figure 16 Processing times for path-rows 181036, 182031, 182032, 182033, 182034and 182035 from 2004 up to 2012.

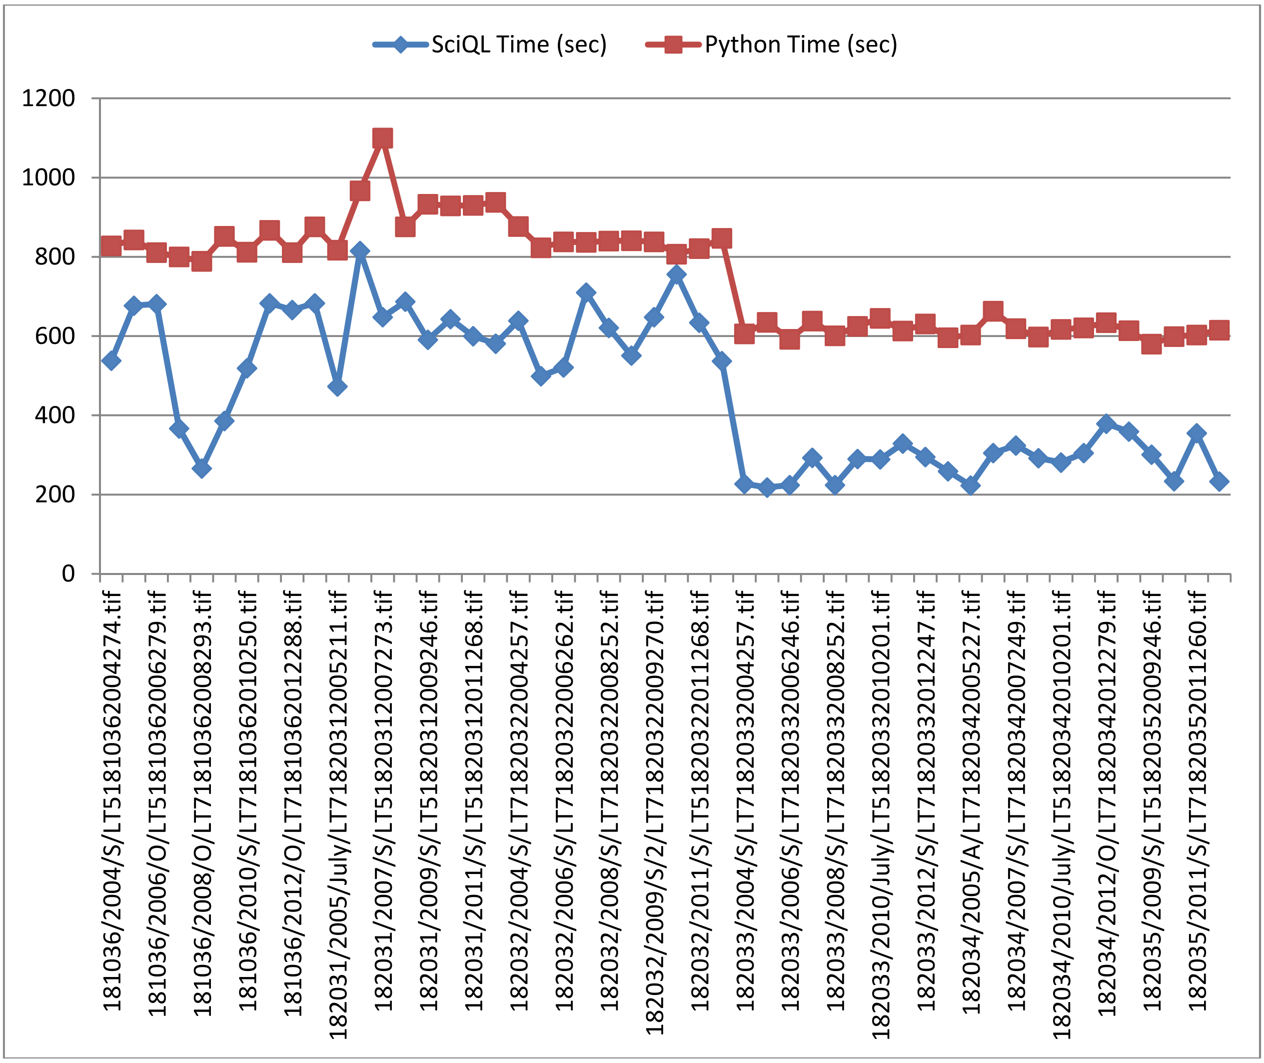

Figure 17 Processing times for path-rows 182036, 183031, 183032, 183033, 183034and 183035 from 2004 up to 2012.

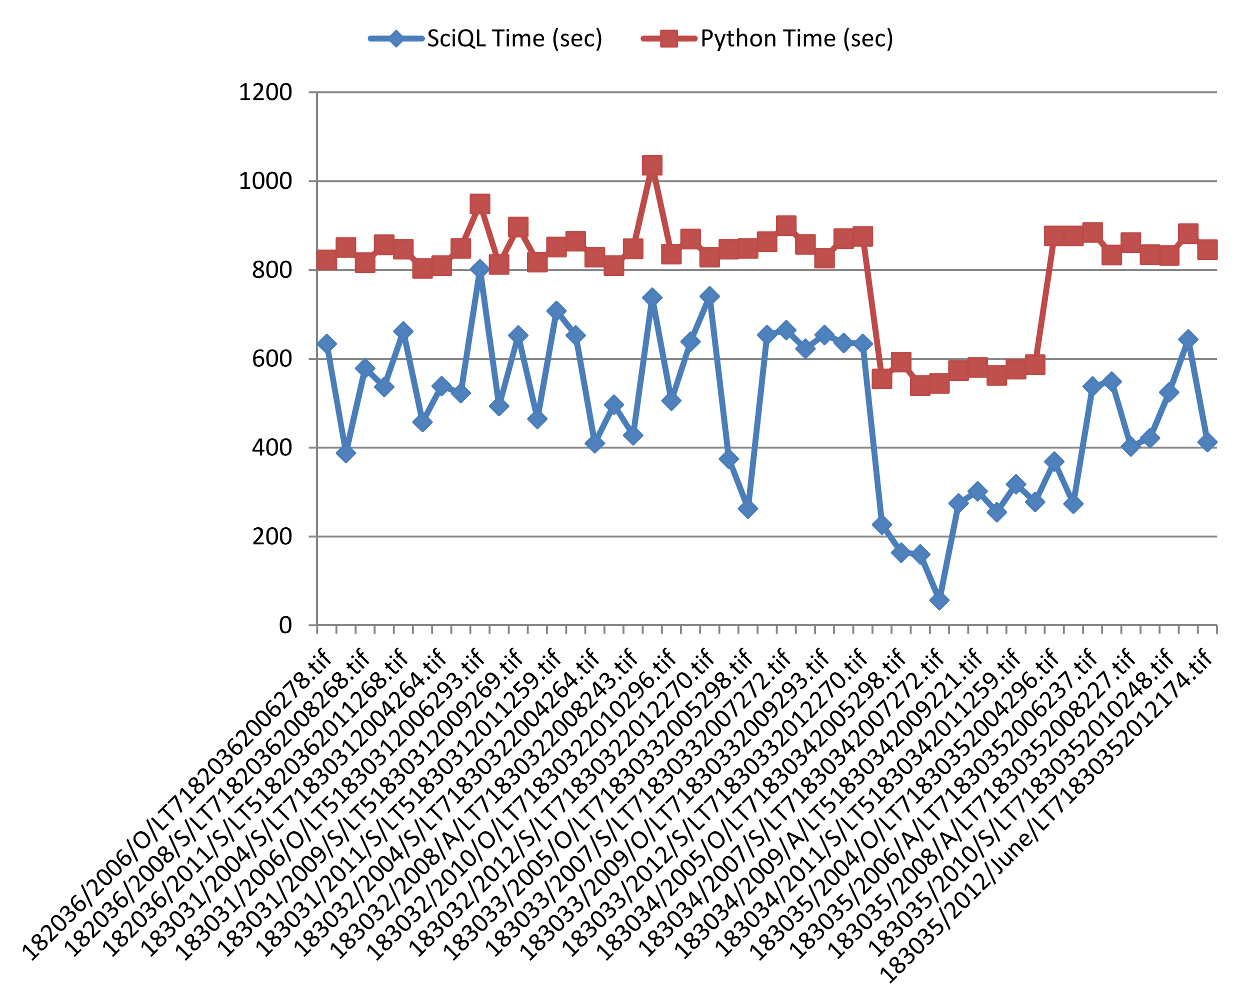

Figure 18 Processing times for path-rows 184031, 184032, 184033, 184034, 185032and 185033 from 2004 up to 2012.

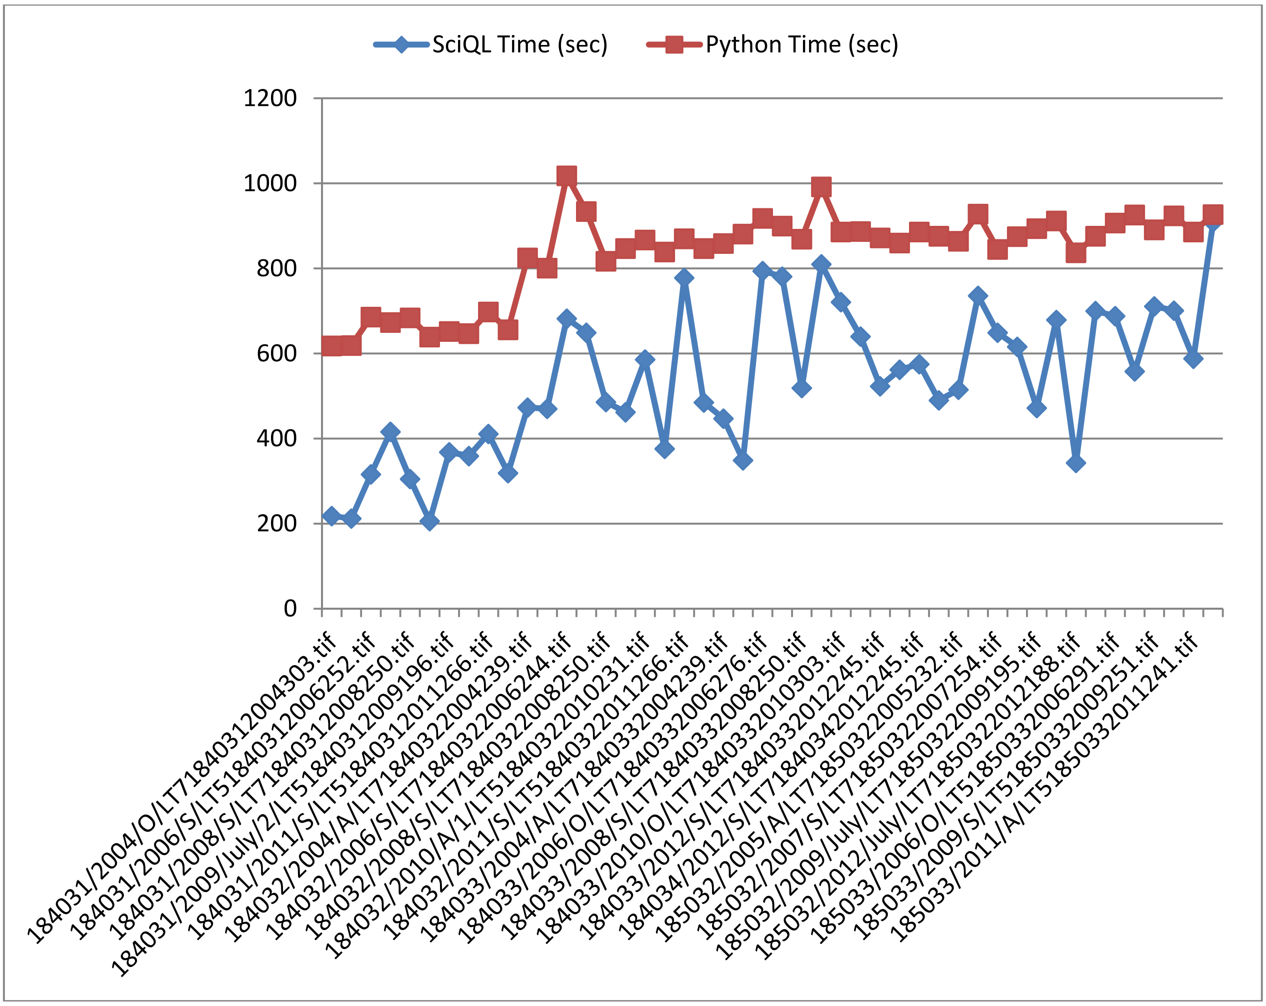

3.3 *Conformity to acceptance criteria*

Acceptance criteria specify the criteria to be fulfilled upon service delivery in order to meet the requirements. From a contractual point of view, they describe the conditions for the decision as to whether the final product fulfills the requirements or not. Acceptance criteria refer to both functional and non-functional requirements. Hence, in the following tables, the focus is to examine what has been accomplished in this second and final version of the fire monitoring use case and which functionalities are still open issues, with respect to the criteria defined in [**5**]. Table 8 inspects the conformity with the functional requirements, while Table 9 with the non-functional requirements. It should be mentioned that the few requirements that were not finally met were rather optimistic definitions which opted for a fully operational implementation. Since TELEIOS is a RTD project, the operational functionalities and 'look and feel' aesthetics that one would expect for an end-to-end service should not all be ready. However, the delivered system has showcased its potential for the deployment of such operational platform, was tested by the user community and received positive feedback.

Table 8. Conformity with the acceptance criteria set for the functional requirements

No	Conditions	Reference to Use Cases	Acceptance criteria	What has been accomplished / open issues
1	Access to the back-end of the system	TELEIOS UC2 SERVICE PROVIDER1, 2, 3,4 and 5, TELEIOS UC2 EOSCIENTIST1	The system backend should be easily and securely accessible	Back-end comprises of independent components (MonetDB, Strabon, front-end), and each one is easily & securely accessible.
2	Ability to modify pre-processing modules	TELEIOS UC2 SERVICE PROVIDER1	Indicatively using a Graphical User Interface and/or a scripting language	Done, through the SciQL based processing chain.
3	Ability to select/specify Look Up Table	TELEIOS UC2 SERVICE PROVIDER1, 4 and 5 and TELEIOS UC2 EOSCIENTIST1	Indicatively using a Graphical User Interface and/or a scripting language	Done, through the SciQL based processing chain and estimation of dynamic classification thresholds
4	Ability to initiate the LUT training process	TELEIOS UC2 EOSCIENTIST1	Indicatively using a Graphical User Interface and/or a scripting language	Done on a one time single step for the definition of the starting values of the static thresholds
5	Ability to create high level queries for refining/adding value to the hot spot and burnt area products	TELEIOS UC2 SERVICE PROVIDER2 and 3	Appropriate querying language and data model. Thematic accuracy as described in the respective non-functional requirement	Done, using stSPARQL queries [**1**]. A graphical query builder is pending but will be delivered in M36.
6	Ability to select projection system for the final products	TELEIOS UC2 SERVICE PROVIDER3 and 5	At least compatibility with UTM and EGSA87 projection systems	At the moment this is done manually inside the SciQL processing chain and not from the TELEIOS interface. Strabon also supports all Coordinate Reference systems.
7	Ability to upload data in the DB	TELEIOS UC2 INTERMEDIATE SERVICEPROVI DER2	> 100 TB storage capacity, and > 99.99% data storage availability	The database has not been stressed to show scalability. This is now achieved via the BSM application (Section 2.2).
8	Ability to define mode of operation between offline or real-time	TELEIOS UC2 SERVICE PROVIDER4 and 5	Indicatively using a Graphical User Interface and/or a scripting language	The functionality has been fully developed (GUI) in the TELEIOS infrastructure
9	Ability to Start/Stop the system	TELEIOS UC2 SERVICE PROVIDER4	Indicatively using a Graphical User Interface and/or a scripting language	The functionality has been fully developed (GUI) in the TELEIOS infrastructure

10	Accessibility of the front-end interface	TELEIOS UC2 ENDUSER1 and 2, TELEIOS UC2 INTERMEDIATE SERVICEPROVIDER1 and 2	Preferably using most common web browsers (e.g. Internet Explorer, Mozilla Firefox, Opera, Google Chrome)	The functionality has been fully developed (GUI) in the TELEIOS infrastructure
11	Functionalities of the web-based interface	TELEIOS UC2 ENDUSER1, TELEIOS UC2 INTERMEDIATE SERVICEPROVIDER1	Availability of all sub-processes described in the respective Use Cases	Most functionalities have been completed / Functionalities not delivered: integration of auxiliary GIS layers, product export capabilities, e-mail notification service, interactive flag updates, GIS data upload by the user.
12	Email notification service	TELEIOS UC2 ENDUSER2	Standard email format with project logos and kml and/or vector data as attachments	Not delivered as the online system is more than enough to serve the initial purpose for dissemination.
13	Administrative privileges	TELEIOS UC2 ADMIN, TELEIOS UC2 ADMIN2	Availability of all sub-processes described in the respective Use Case	A super user with administrative privileges has been created.

Table 9. Conformity with the acceptance criteria set for the non-functional requirements

No	Conditions	Reference to requirements	Acceptance criteria	What has been accomplished / what is due in the V2
1	Automatic generation of hot spots time limit	UC2.PR.01	Less than five minutes	Accomplished [**1**]
2	Complex query processing time	UC2.PR.03	Prompt system response	Accomplished [**1**]
3	Handling of terabytes of data without performance losses	UC2.PR.02	Prompt retrieval of raw data and related by-products	Accomplished for the real-time fire monitoring use case, however scalability surfaces through the BSM application
4	High availability/robustness of the system	UC2.RR.01	99.99% system availability	Tested (i) in the framework of the 3rd TELEIOS user workshop (Section 4.1.3), and (ii) consortium stress tests of the online system. The targeted availability has been achieved
5	High thematic accuracy	UC2.RR.02	Preferable false alarm rate < 15% Preferable omission errors < 15%	Accomplished (see Section 3.1)
6	High positional accuracy	UC2.RR.03	On a sub-pixel basis hot spot products = 2 km burnt area products = 500 m	Accomplished for the real-time fire monitoring use case via the application of an accurate resampling polynomial
7	Capability to export to vector	UC2.IR.01	compatibility with OGC standards (indicatively GML 3.1 SF0)	Not delivered

	data			
8	Capability to export to geo-referenced raster data	UC2.IR.02	Preferably to img and geotiff formats	Not delivered
9	Capability to export an interactive map view to an image file	UC2.IR.04	Indicatively to jpg format	Not delivered
10	Export to Google Earth compatible data	UC2.IR.03	Export to kml format	Not delivered
11	Uploading and viewing metadata	UC2.IR.06	Metadata compatibility with ISO 19115 and/or INSPIRE standards	Not delivered
12	Consistency of the information support	UC2.IR.07	Consistency between map and legend symbols and compatibility between the geographic coordinate/projection systems of the various sources	Done, no inconsistencies are applicable
13	Content clarity of the front-end interface	UC2.IR.08	Indicatively, visual identification of 1) an overview map, 2) coordinate Graticules/Grid, 3) cartographic scale, 4) interpretation text, 5) project logos, 6) reference of the information sources, 7) acknowledgments, 8) constraints related to access, use and information sharing and 9) liability information	Not delivered
14	Security of the products in the DB	UC2.SR.01	Typical security options incorporated in DBs	Done
15	Efficient front-end authorization mechanism	UC2.SR.02	Distinct user privileges	Done, users with diverse privileges can be created
16	Documentation completeness	UC2.DR.01, 02, 03 and 04	Existence of installation, administrator, end user and database schema manuals	Done. More information can be found in [**12**] and [**3**].

4 User feedback

4.1.1 1st TELEIOS user workshop

A broad spectrum of users attended the 1st TELEIOS user workshop held in Frascati, including stakeholders from the EO and IT communities, end-users (civil protection agencies, environmental and governmental agencies in Greece like HMOA/DGF, FRI/NAGREF, HCBW and GSCP) [**6**]. The various user roles of a fire monitoring system were identified:

1. END_USER, using fire monitoring products on an operational basis,

2. EO_SCIENTIST consisting of research groups in the field of remote sensing,
3. SERVICE_PROVIDER, that delivers added-value products from raw EO imagery,
4. INTERMEDIATE_SERVICE_PROVIDER, that delivers added value products on the basis of other ready products generated by third parties,
5. ADMINISTRATOR, who has full access to the system and may change configuration data.

The stories of users, coming from their daily experience were collected, and by grouping these user stories the corresponding use cases per user role were derived. The latter served as a complete set of functional requirements that the fire monitoring system should satisfy. Additionally, several non-functional requirements were determined, associated with performance, reliability, interface, security, storage and persistence, and documentation issues. Finally, acceptance criteria that need to be fulfilled upon service delivery were specified. These criteria were presented in the previous section.

4.1.2 2nd TELEIOS user workshop

The workshop, which was held in Darmstadt in May 2012, gave the opportunity to the TELEIOS partners to highlight the advancements accomplished in the real-time fire monitoring application using database and semantic technologies. The scenarios presented included the functionalities that were integrated so far, aiming at improving the capability to extract information from single images and from time-series of images, thus supporting the work of real-time decision makers and EO scientists:

1. Real-time fire monitoring with simulated arrival of new imagery.
2. Search for raw satellite data.
3. Offline processing of raw data using MonetDB, and ability to select different processing chains.
4. Search for hotspot products.
5. Refinement of hotspot products with semantic queries.
6. Enrich the automatic generation of fire maps with relevant geo-information.

While the version of the system that was presented was in a development phase, overall the fire monitoring application was very well received. The main issues, as expressed by the users, which should be addressed in more detail referred to the data access from distributed places, scalability to deal with larger areas, more robust integration of the different system modules, and licensing policies for making data and products accessible for everyone.
More details on the outcomes of the 2nd TELEIOS user workshop can be found in D8.2.1 [**7**].

4.1.3 3rd TELEIOS user workshop

A thorough analysis of the feedback collected during the 3rd TELEIOS user workshop, held in Chania on 12th of June 2013, is included in D8.2.2 “An evaluation of the TELEIOS infrastructure (Final version) by the TELEIOS user community” [**21**]. Some comments are included in this section, however, in order to provide the users’ view point in the context of the real-time fire monitoring use case.

To serve the needs of the user workshop we set up a local server with the following technical characteristics:

- CPU: Intel i7-3770K (3.5 GHz, 8 MB L3 cache)

- HD: 2 x 2TB (7.200 rpm, 64 MB cache)
- RAM: 32 GB

In this machine we installed 4 Virtual Machines (VM), where each VM was configured to have 2 cores and 6.5 GB of RAM. The TELEIOS system was installed to a each of the four VMs, in order (i) to minimize dependencies on a robust internet connection, and (ii) avoid having a setting where different users update simultaneously the database with conflicting statements. Therefore each distinct user, from a total of 16 users, was handed a clear scenario he/she had to carry out. Indicatively, a specific scenario that was actually implemented by several users follows:

Group 1.

Login as a super-user (credentials were given separately).

Phase A

Step A.1: Go to the 'Offline Functionality' tab. Search for hotspots from 25/08/2007 23:00:00 until 25/08/2007 23:59:00.
Step A.2: Play with the various GIS functionalities (select/invert, pan/zoom, view hotspot details, etc.) on the bottom right corner.
Step A.3: Toggle between raw and refined hotspot for timestamp 25/08/2007 23:15:00 and observe the differences.
Step A.4: Clear search.

Phase B

Step B.1: New search for hotspots: from 22/08/2010 04:40:00 – 22/08/2010 05:35:00.
Step B.2: Notice whether some hotspot(s) are missing for your target timestamp: *22/08/2010 05:10:00*, with respect to older timestamps.
Step B.3. Apply the 'Delete in Sea Refinement' and 'Partial Sea Hotspot refinement' for the timestamps that are 30 minutes older than your target timestamp.
Step B.4. Apply all available refinements for your target timestamp.

Phase C

Step C.1: New search for hotspots from 24/08/2011 10:31:00 – 24/08/2011 10:45:00.
Step C2: Move to the 'Offline Processing' tab and search for data from 24/08/2011 10:31:00 – 24/08/2011 10:45:00.
Step C3: Move to the 'Import New Data' tab and go to folder /2011 and import the data (both 039 and 108 bands) that correspond to the previously searched time range.
Step C.4: Redo step C.2, select all items, choose 'Processor Version' -> 'Dynamic Threshold' and press 'RUN'. We are automatically transferred to the 'Offline Scheduler' tab. Track progress by pressing the 'Refresh Scheduler Grid' button.
Step C.5: Redo step C.1.
Step C.6: Apply a 'Time Persistence Refinement' operation. Clear search and redo step C.1.
Step C.7. Redo steps C.1 – C.6 for a new time period: 24/08/2011 16:01:00 – 24/08/2011 16:15:00. Try to use both 'Static Thresholds' and 'Dynamic Thresholds' to evaluate the improvement achieved by using a more advanced processing chain. Execute the processing chain for a third time activating this time the 'Add Refinement Chain' option.

The above script basically provides, in three distinct phases, a flavor for all available functionalities and improvements that have been achieved in the framework of TELEIOS. Starting from the basic 'search for hotspots' functionality (Phase A), moving to refining certain products using semantic queries on top of Strabon (Phase B) and applying the entire process (Phase C) that includes importing raw MSG/SEVIRI data, triggering the SciQL processing chain (with the ability to choose among different versions of the classification module), and applying custom refinement queries.

In general the feedback collected, which will be elaborated and analyzed in [**22**], was very encouraging: most users found the applications very useful, specifically when it concerned stakeholders that need fire monitoring products as part of their daily work practice (e.g. Greek Ministry of Environment, Energy and Climate Change, Italian civil defense agency, or foresters in local administrative units). Comments regarding the underlying technologies (MonetDB and Strabon) were minimal, while most criticism referred to 'look and feel' issues of the front-end interface. An operational web-based tool for fire monitoring was not however the aim when developing the platform, but it was merely to ensure the coherent and integrated operation of different technologies that were developed in the framework of TELEIOS. Therefore the weight was deliberately set upon increased system availability, robust and timely processing, with some basic web-GIS functionalities. However for reasons of completeness some related comments made by the users regarding the interface usability are included:

- Map element should be more interactive with the mouse
- Not possible to change the order of the processes or to cancel a process in one of the operations
- The way of entering dates could be improved, e.g. by using a time slider (this was an issue raised by several users)
- Map view: hard to find the legend for the map layer, more support in the localisation through boundaries would be nice to have, screen resolution can lead to loss of part of the interface
- Flash application does not run on mobile devices
- Data import: you have to double click in order to be able to import files – it should be possible to load the whole directory
- Map view looks a bit outdated compared to the current state of web mapping applications
- Need to export the fire pixel classification to the form of vector- polygon not only raster, need to send in real time the observed fire pixels in the form of polygon to the control room for further analysis.

5 Lessons learnt and future plans

In the framework of TELEIOS two different scientific communities came closer, broadly speaking the remote sensing and the information technology communities. Array database and semantic technologies have come to tackle a real need in Earth Observation, the effective processing, management and exploitation of big data volumes that are incorporated in sophisticated algorithms and workflows. These technologies were represented in TELEIOS by CWI and NKUA partners, while NOA and ACS were the partners responsible for integrating them and rendering them as pre-operational services.

In WP7, two service chains were successfully deployed using tools that were developed in TELEIOS, namely the real-time fire monitoring and the burnt scar mapping services. This entailed the close collaboration between the partners by establishing a robust feedback mechanism. It can be argued that in the aftermath of TELEIOS, NOA has appreciated the potential benefits of using cutting-edge IT technologies, and NKUA & CWI have acquired the necessary criteria to assess in which scientific scenarios and problems they can offer useful solutions to EO scientists.

Since this deliverable is written from the perspective of the remote sensing scientists, the intense interaction with the TELEIOS technologies provided some lessons that will be taken into account in the NOA activities beyond TELEIOS.

With respect to array database tools, we were able to identify the following lessons learnt:

- SciQL implementation of existing processing chains written in other programming languages (e.g. C or SciPython) leads to concise queries and scripts. We can express common EO operations easily using the purpose-build SciQL instead of using a lengthy C program. However, as always, it is necessary to devote some time in learning SciQL, although it is intuitive and the learning curve is rather steep.
- The code length in our processing chains for fire monitoring and burnt scar mapping is concisely written in SciQL, making debugging easier and allowing potential functionalities extensions easier to add and track. Hence, rapid prototyping is feasible and new processing modules can be created without the need to recompile everything.
- Not all mathematical operations have been developed, therefore in this early stage new applications require effort from CWI for the development of new functionalities. This is to be expected however, since this is work in progress.
- The absence of an in-depth user's manual for SciQL makes independent work difficult, and interdependence with CWI developers is implied.
- The Data Vault functionality has proven its usefulness, as it hides completely the details/complexity of the raw satellite data format.
- The use of higher-level languages like SciQL, allowed NOA to stop worrying about how to store and manage data and just focus on the actual processing.
- Additionally, as shown specifically in the BSM application, we were able to easily handle massive data loads due to advanced processing capabilities of column-store MonetDB.
- In terms of the processing times, as seen in Section 3.2, processing on top of MonetDB is quite fast and poses no restrictions. Algorithm execution is optimized by the DBMS's query optimizer, and there is room of improvement to this end, speed-wise.

With respect to semantic technologies, we were able to identify the following lessons learnt:

- The value of applying semantic queries for the thematic refinement of the hotspot products has been appreciated both by NOA and the user community NOA serves. Via the implementation of spatio-temporal queries expressed in stSPARQL, we were able to identify and eliminate hotspots occurring in the sea, decide for hotspots that are partly located in non-consistent underlying land use, and attribute a variable confidence level according to the spatio-temporal persistence of a hotspot.
- The processing speed is satisfactory for the intended applications. It should be mentioned, however, that there is some effort involved for posing optimum SPARQL queries. Achieving the good processing times that are presented herein was not a straightforward task.

- It was shown, and it is envisaged, that rapid mapping applications can be easily deployed using semantic technologies with distributed data. We can use other Linked Data together with fire products (e.g. hostpots or burnt areas) to further enhance our products' value and create thematic maps that can be used by stakeholders. Greek government data (geodata.gov.gr), Administrative Geography of Greece, Open Street Map, Wikipedia, Gazeteers (e.g. Geonames) can be incorporated in a couple of simple steps.
- New value can be extracted from existing data sources, by expressing semantic queries like "Find hotspots within 2 km from a major archeological site" or "Find high value forested properties near the sea that were burnt more than once and have been subsequently built".
- Strabon has reached a level of maturity that can accommodate real-time applications, ensuring the seamless operation of end-to-end services.

MonetDB and SciQL, along with Stabon and SPARQL have proven their value in the framework of this Use Case, in WP7 of TELEIOS. The next step would be to use them extensively within the various research activities of NOA. NOA is set to become a regional EO center for monitoring natural disasters. To this end, NOA has been recently awarded with a capacity building project, BEYOND[11]. BEYOND aims to maintain and expand the existing state-of-the-art interdisciplinary research potential, by Building a Centre of Excellence for Earth Observation based monitoring of Natural Disasters in south-eastern Europe, with a prospect to increase it access range to the wider Mediterranean region through the integrated cooperation with twin organizations. Through BEYOND it will be possible to:

- Set up innovative integrated observational solutions that will allow a multitude of monitoring networks (space-borne and ground-based) to operate in a complementary, unified and coordinated manner.
- Create archives and databases of long series of observations and derived higher level products, and
- Make these observations and products available for exploitation with the involvement of stakeholders, scientists and/or institutional users, applicable for down-streaming to their specific needs.

Additionally, NOA is close to reaching an agreement with European Space Agency for setting up a mirror site in its premises to archive and disseminate in real-time data from the Sentinel family of satellites, adopting an open data strategy.

Therefore BEYOND, in conjunction with the ESA mirror site, are envisaged to offer unique opportunities to NOA for creating new and scalable applications based on big data. The signal and EO data processing know-how at NOA has been long established. Through participation in TELEIOS and familiarizing with semantic and array database technologies, NOA is ready to offer hereafter leading-edge solutions via the effective exploitation of the anticipated big data flow.

[11] http://www.beyond-eocenter.eu/

6 List of Abbreviations

AOI	Area Of Interest
BSM	Burn Scar Mapping
CLC	Corine Land Cover
CP	Control Point
DEM	Digital Elevation Model
DoW	Description of Work
EO	Earth Observation
ERS	Emergency Response Service
FIRMS	Fire Information for Resource Management System
FMM	Fire Monitoring at Middle resolution
FRI/NAGREF	Forest Research Institute of Agricultural Research Foundation
GIS	Geographic Information System/Services
GMES	Global Monitoring for Environment and Security
GSCP	General Secretariat for Civil Protection
HCBW	Hellenic Centre of Biotopes and Wetlands
HRIT	High Rate Information Transmission
HMOA/DGF	Hellenic Ministry Of Agriculture/Directorate of Forests and Natural Resources Development and Protection
IT	Information Technology
JRC	Joint Research Center
LGD	LinkedGeoData
LOD	Linked Open Data
LRIT	Low Rate Information Transmission
LUT	Look Up Table
MODIS	Moderate-Resolution Imaging Spectroradiometer
MSG	Meteosat Second Generation
NOA	National Observatory of Athens
OGC	Open Geospatial Consortium
ONF	Office National des Forêts of Corsica
OSM	Open Street Map
RDF	Resource Description Framework
SAFER	Services and Applications For Emergency Response
SciQL	Scientific Query Language
SEVIRI	Spinning Enhanced Visible and InfraRed Imager
SP	Service Provider
SPOT	Système Probatoire d'Observation de la Terre
SQL	Structured Query Language
VM	Virtual Machine

7 References

1 TELEIOS. *D7.3 'Implementing the real-time fire monitoring application – Phase II'*. 2013.

2 TELEIOS. *Description of Work*.

3 TELEIOS. *D1.4 'The TELEIOS infrastructure - version II'*.

4 TELEIOS. *D3.1 'KDD concepts and methods proposal: report & design recommendations'*.

5 TELEIOS. *D7.1 'Requirements specification for the real-time fire monitoring application'*.

6 TELEIOS. *D8.1 'Requirements specification of the TELEIOS user community*.

7 TELEIOS. *D8.2.1 'An evaluation of the TELEIOS infrastructure (Version 1) by the TELEIOS user community'*.

8 Sifakis, N., Iossifidis, C., Kontoes, C., and Keramitsoglou, I. Wildfire Detection and Tracking over Greece Using MSG-SEVIRI Satellite Data. *Remote Sensing*, 3 (2011), 534-538.

9 EUMETSAT. *Active fire monitoring with msg algorithm - theoretical basis document*. Darmstadt, 2007.

10 TELEIOS. *D5.1 'An implementation of ad-hoc array queries on top of MonetDB.*

11 TELEIOS. *D4.1 'An implementation of a temporal and spatial extention of RDF and SPARQL on top of MonetDB - phase I'*.

12 TELEIOS. *D1.3 'The TELEIOS infrastructure - version I'*.

13 Carvalheiro, L.C., Bernardo, S.O., Orgaz, M.D.M., and Yamazaki, Y. Short communication: Forest fires mapping and monitoring of current and past forest fire activity from meteosat second generation data. *Environ. Model. Softw.*, 25 (2010), 1909-1914.

14 Calle, A., Casanova, J., and Romo, A. Fire detection and monitoring using MSG Spinning Enhanced Visible and Infrared Imager (SEVIRI) data. *Journal of Geophysical Research*, 111, G4 (2006).

15 Calle, A., Gonzalez-Alonso, F., and Merino de Miguel, S. Validation of active forest fires detected by msg-seviri by means of modis hot spots and awifs images. *International Journal of Remote Sensing*, 29 (2008), 3407-3415.

16 Kontoes, C., Papoutsis, I., Michail, D. et al. Wildfire monitoring via the integration of remote sensing with innovative information technologies. (Vienna 2012), EGU2012.

17 TELEIOS. *D2.1 'A data model and query language for an extension of RDF with time and space'*.

18 Kontoes, C., Poilvé, H., Florsch, G., Keramitsoglou, I., and Paralikidis, S. A

comparative analysis of a fixed threshold vs. a classification tree appoach for operational burn scar detection and mapping. *International Journal of Applied Earth Observation and Geoinformation*, 11 (2009), 299-316.

19 TELEIOS. *D5.3 'An Evaluation of Ad-hoc queries on arrays in MonetDB'*. 2013.

20 SAFER. *D21000-3'Portfolio of thematic and technologic innovation services within the project and impact on operational architecture'*. 2010.

21 TELEIOS. *D8.2.2 'An evaluation of the TELEIOS infrastructure (Final version) by the TELEIOS user community'*. 2013.

22 TELEIOS. *D8.2.2 'An evaluation of the TELEIOS infrastructure (Final version) by the TELEIOS user community'*.

23 Kontoes, C., Keramitsoglou, I., Papoutsis, I. et al. Operational Wildfire Monitoring and Disaster Management Support Using State-of-the-art EO and Information Technologies. (Beijing 2012), Second International Workshop on Earth Observation and Remote Sensing Applications.

24 Kontoes, C. Operational land cover change detection using change vector analysis. *Journal of Remote Sensing*, 29, 16 (2008), 4757-4779.

25 Koubarakis, M., Kyzirakos, K., Karpathiotakis, M. et al. TELEIOS: A Database-Powered Virtual Earth Observatory. (Istanbul 2012), 38th International Conference on Very Large Databases (VLDB). Demo paper.

26 Nikolaou, C., Dogani, K., Kyzirakos, K., and Koubarakis, M. Sextant: Browsing and Mapping the Ocean of Linked Geospatial Data. (Montpellier, France 2013), 10th Extended Semantic Web Conference (ESWC 2013).

27 Kyzirakos, K., Karpathiotakis, M., Garbis, G. et al. Real Time Fire Monitoring Using Semantic Web and Linked Data Technologies. (Boston, USA 2012), 11th International Semantic Web Conference (ISWC 2012).

28 Bereta, K., Smeros, P., and Koubarakis, M. Representing and querying the valid time of triples for Linked Geospatial data. (Monpellier, France 2013), 10th Extended Semantic Web Conference (ESWC 13).

29 TELEIOS. *D4.2 'An implementation of a temporal and spatial extention of RDF and SPARQL on top of MonetDB - Phase II'*.

30 Koubarakis, M., Kontoes, C., Manegold, S. et al. Real-time Wildfire Monitoring Using Scientific Database and Linked Data Technologies. (Genoa, Italy 2013), 16th International Conference on Extending Database Technology (EBTD 2013).

8 List of Figures

9 List of Tables