GEOSS AND THE SENSOR WEB

GEOSS Sensor Web Workshop Report

June 2008

ingo.simonis@geospatialresearch.de
echterhoff@uni-muenster.de

GEOSS AND THE SENSOR WEB

GEOSS Sensor Web Workshop Report

Ingo Simonis - Meraka Institute, CSIR, South Africa
Johannes Echterhoff - IfGI, University of Muenster, Germany
June 2008

Abstract

On May 15 and 16, 2008, more than 40 researchers, members of the GEO Secretariat as well as representatives of the European Commission met at the GEO/WMO premises in Geneva, Switzerland, to discuss ongoing activities, concepts, missing bits and pieces as well as new approaches in the context of Sensor Web. This was the second workshop organized by GEO Task DA-07-04, led by the Meraka Institute in South Africa and greatly supported by the GEO Secretariat. The workshop was further co-organized by the European Commission, Directorate-General Information Society and Media (DG INFSO) with the support of two integrated projects SANY [1] and OSIRIS [2]. This report provides an annotated in-depth summary of the workshop discussions and results and provides extensive background information.

The Sensor Web: An active Sensing Macro Instrument

Introduction

GEOSS has the objective to continuously monitor the state of the earth in order to increase knowledge and understanding of our planet and its processes. Being a system of system, GEOSS has to master the challenge of integrating heterogeneous systems across institutional and political boundaries. Timely delivery of earth observation data is a key aspect in identifying potential natural and human threats, such as tornados, tsunamis, wild fires, or algae blooms.

Data from in-situ or remote sensing devices form the basis for analyzing gradual processes, such as increasing drought, water shortages, or rising sea levels. The Sensor Web [3] presents a paradigm in which the Internet is evolving into an active sensing macro instrument. Based on internationally adopted standards, the Sensor Web accomplishes the necessary requirements to ensure interoperability among its various components. Sensor Web software components and services work interactively in arbitrary application domains.

The aim of the Sensor Web Workshop was to

- discuss scenarios and use cases to improve interoperability between sensor data providers and consumers
- identify standards and best practices that should be part of the GEOSS Information Infrastructure Architecture
- learn about the Sensor Web concept, its potential and role within GEOSS
- explore the capabilities of a strong Sensor Web within GEOSS
- discuss about the various approaches, applications, and technologies used within sensor networks.

Workshop Course

The workshop started with a general introduction to GEOSS task DA-07-04 and its particular role in GEOSS. The ensuing discussion showed that the Sensor Web simultaneously is a system

- still under development that needs additional research, and
- that is already put into action in various application domains.

The remaining day was dedicated to presentations about the newest developments in international Sensor Web related research and industrial projects. The great variety of Sensor Web R&D and Sensor Web implementation and usage was reflected by 25 presentations given by workshop participants; all to be found online at OGC Network [5]. The workshop participants split into three working groups on the second day to discuss a number of research questions and to develop the task sheet of GEOSS task DA-07-04 further.

Use of Standards

The importance of standards was emphasized in almost all presentations and discussion threads. The standards developed by the Sensor Web Enablement (SWE) initiative of the Open Geospatial Consortium (OGC) [6] take on a leading role in this context. The standards "Sensor Observation Service" to access sensor data, "Sensor Planning Service" to task sensors, "Sensor Alert Service" to subscribe to events, and "Web Notification Service" to deliver messages on a variety of communication protocols are used through out most of the current projects with sensors involved.

The standards rely on specific models and encodings to describe sensors and sensor systems (SensorML) as well as observations (O&M). Interestingly enough, most implementations are based on Open Source Software that is provided by the Open Source initiative 52°North under the GPL License [4]. A number of research and industrial projects contribute to this initiative and help to facilitate set-up and enhancement of Sensor Web applications. Commercial solutions are available as well. The Sensor Bay project [9] used Compusult's Web Enterprise Suite (WES) to set up a SWE compliant sensor network in remarkably short time and budget.

Though widely used in a very successful manner, SWE standards still show some limitations. Without claiming to be exhaustive, we will discuss a number of items in the following that will have to be addressed in the future.

Sensor Registries, Dictionaries & the Role of Semantics

Discovery of data sets in general and in GEOSS in particular depends on a specific set of registries and dictionaries that allow data and service discovery. Currently, the initial operating capability of GEOSS consists of three types of registries, none of them reflecting the specific requirements of the Sensor Web. The situation in the context of dictionaries is even worse. The lack of GEOSS Sensor dictionaries prevents participants in GEOSS from using homogenous terms that could then be used for sensor, data, or service discovery. OGC Sensor Web Enablement is of little help in the regard either: A common registry is not yet available. Work in this research area is progressing, however, to be applicable in GEOSS, a registry has to be developed that handles specific aspects of the Sensor Web, including a key characteristic: Information coming from sensors as well as sensors themselves are often highly dynamic. In detail, the following aspects need further research:

- SWE standards address the provision of sensor data through Web services. A registry has to be able to harvest the information provided by individual services. In addition to this harvesting process, the registry has to ensure that information about individual sensor instances is captured properly. The links between sensors and the services that encapsulate them should be discoverable, taking into account that a single sensor might be accessible through a set of several service instances. Also the time-dependency of such relationships (e.g. for moving sensors) was considered by the workshop participants to be useful information that should be covered by registries.

- Talking about highly dynamic sensor networks (which may involve ad-hoc wireless sensor networks or other self organizing structures), the need for sophisticated discovery mechanisms (like ad-hoc / distributed search algorithms) was emphasized. These mechanisms should be able to cope with metadata changing in time (e.g. calibration information about sensors and sensors "moving" around in a service environment) and with time dependent availability of sensors. The dynamic aspect is important for service instances as well, as a service installed on a sensor node will go down together with the sensor itself when battery power is exhausted.

- A sensor web registry should be able to interpret and support the data encodings commonly used in the Sensor Web. Specifically SensorML and O&M should be supported as standard encodings for describing sensors and observations.

- The registry has to support actively those operations necessary to harvest metadata from other services. Work in this area has to facilitate harmonization of operations commonly used for retrieving information (e.g. about sensors and observations) across the various service specifications.

- A big issue is the lack of semantically rich discovery mechanisms in SWE services based infrastructures. Search algorithms facilitating semantically enhanced queries from users would be of great benefit for retrieving useful information out of Sensor Web registries and services. At the moment, related concepts, subgroups of sensor types, or other dependencies can hardly be explored. Eventually, the integration of domain ontologies (e.g. the SWEET ontology, see http://sweet.jpl.nasa.gov/ontology/), semantic queries and semantic transformations in Sensor Web infrastructure have to be addressed.

Dynamic Sensor Networks with Moving Sensors

Roaming of sensors inside a sensor network has not been thoroughly considered by SWE activities so far. Consequently, some deficits exist that should be handled by enhancing the relevant standards. The Sensor Observation Service (SOS) was discussed with respect to this topic. A problem of the current service interface is that features of interest (providing the spatial relationship of observations) are provided in a plain list inside the Capabilities document of an SOS instance. When moving sensors are registered at a SOS, a new feature will be created whenever a new position is reported, which leads to a dramatically increase of the size of the capabilities response. Potential solutions to this problem include paging mechanisms, specialized operations for retrieving features of interest by region, and associating the measurements of moving sensors to domain features while considering the spatial information part of the observation data.

Enhancements of the existing SWE interfaces should focus on the aspects

- discovery, exploration, and retrieval of dynamic data, and

- insertion and updates of time dependent data

For example, retrieving a SensorML description via the “DescribeSensor” operation of a SOS based upon a given time frame is not yet possible – only the whole document can be accessed.

Metadata Information in SensorML

Metadata about sensors, platforms, and sensor data is a critical aspect of Sensor Web infrastructures. Following the developments in the context of the European directive INSPIRE (http://www.ec-gis.org/inspire/), Metadata according to ISO standards 19115 and 19119 shall be supported by SWE standards. Though first metadata profiles have been discussed (good examples exist in particular in the remote sensing community, see discussion of Cal/Val portal below), we are still far from having mature metadata models or even policies and procedures that could be applied to a harmonized development of metadata across application domains.

SensorML documents are currently developed from scratch or by modifying existing examples, rather than by guidelines and best practices. The results are SensorML instances that are barely interoperable. Work needs to be done to identify best practices and relevant standards as well as to create guidelines for defining and incorporating quality information.

Several projects are dealing with these topics, focusing on special aspects. Two projects are listed here for reference:

- The CalVal portal (http://www.brockmann-consult.de/CalValPortal/welcome.do) is providing information and data for Calibration (Cal) and Validation (Val) of Earth Observation (EO) data, thereby offering the basis for further work in this area.
- The INTAMAP project [10] produced a XML Schema based encoding called UncertML (http://www.intamap.org/uncertml/uncertml.php), which deals with description of uncertainty information and might qualify as a language to overcome some of the limitations of SensorML and O&M (here, only self-defined fields are currently used to express quality and uncertainty information). This is a particularly important aspect of sensor networks, as sensor data is rarely used as an end product. Usually, higher-level information is derived out of low-level sensor data. UncertML does not reflect standard practices like the ISO "Guide to the Expression of Uncertainty in Measurement" (GUM).

Verbosity and flexibility of SensorML

The description of all details of sensors may lead to verbose documents. An extreme example is the description of MERIS. MERIS is a programmable, medium-spectral resolution, imaging spectrometer operating in the solar reflective spectral range. With five cameras resulting in 15 bands and two tables for each band of each camera, the full description sums up to 18,000 lines of code. Still, this is an extreme example. Other sensor descriptions merely reach a hundred lines of code and are fairly easy to process and encode. A more severe problem in terms of interoperability is the flexibility SensorML provides. Though a strength of SensorML on the one side, the lack of mandatory elements requires each community to produce their own profile in order to reach the highest level of interoperability. As various communities use single data sets for entirely different purposes, the necessary elements in the description of the sensor and the data processing chain vary. On a more abstract level, the decision to optimize a SensorML document towards technical details, potential usage, or integration issues results in different implementations.

In summary, the creation of universal usable SensorML documents is not practical. Even the separation in little chunks of SensorML code that will be aggregated to meet individual needs is of little help at the moment, as mechanisms to discover and integrate SensorML fragments are still missing. Collecting SensorML examples and best practices would help on this regard to use SensorML in a more efficient way and to support its uptake in various application domains. The Cal/Val portal [8] or the GEOSS portal would be suitable places to host information about SensorML, examples, and best practice reports.

Sensor Supervision

Workshop participants raise the issue of sensor supervision as an important aspect of the Sensor Web. Sensor providers as well as users should be able to access the status of SWE services and the sensors associated with them. This would help to identify unusual sensor behavior and malfunctions as well as to identify current location, battery level, health, parameterization, tasks etc. of a sensor. The information provided could foster timely reaction to situations of interest – not only for end users but also for the whole sensor network or Sensor Web in general, thereby improving the reactivity of the whole application.

Operations (and services implementing them) need to be defined for supporting supervision, basically allowing insertion and retrieval of (probably highly dynamic) sensor status information. Such functionality would also be relevant for sensor discovery, i.e. for SWE registries. For example, clients could search for all sensors currently available in a certain area of interest.

Security

There are several levels at which a Sensor Web infrastructure might need to be secured. Although secured networks are not a general goal of GEOSS, it is obvious that sensor providers will not open their networks to anyone in every case. In addition to access control, the aspects integrity, non-repudiation, and confidentiality will have to be addressed.

The participants presented several use cases and identified the following (non exhaustive list) of requirements:

- Access to certain observations had to be restricted. For example, common users might only be allowed to view radioactivity data that is older than one week (giving authorities time to react upon the latest information).

- Allowing insertion of data into the service environment could be secured to prevent false information and to preserve the integrity of available data.
- When tasking sensors like for example an unmanned aerial vehicle (UAV) or satellite, different users might be granted access to different tasking functionality. Roles could be assigned to users, which would allow access to pre-defined sets of operations.
- Tasking requests to some assets must ensure non-repudiation. During transfer of data, confidentiality has to be ensured. In some cases, integrity of requests and responses has to be guaranteed.

Sharing Experiences and Best Practices

Many research and industrial projects are dealing with in-situ and remote sensor technologies. More and more of those projects implement SWE-services, not always using the same approach for encoding specific data types and describing specific sensor types. Although these infrastructures will be interoperable from a general point of view, actually the integration of data sets from different deployments is an ambitious task.

Thus, as was said earlier, best practices and guidelines need to be established to facilitate interoperability between different infrastructures and within GEOSS. The workshop participants would appreciate to have a common portal (e.g., the GEO portal) where all projects dealing with SWE could be listed. For each project, additional information e.g. about their service architecture and their approach for encoding data (in SensorML and O&M) should be made accessible. Means for providing feedback and discussing SWE in general as well as the listed projects should be provided by the portal. This could be a forum or a Wiki, for example. Based upon feedback and discussions taking place on the portal, further best practices could be established and documented.

Another idea was to gather use cases where SWE has been, will be or could be applied. From these use cases guidance could be provided where and how SWE is suitable for deployment. This information should also be gathered on the common portal.

Future of GEOSS Task DA-07-04

One of the most important outcomes of the workshop was the extension of GEOSS task DA-07-04 to include sensors being deployed under water. DA-07-04 now includes the oceanographic community. The updated task sheet is available at GEOSecretariat's FTP site [7].

Next Workshop

We received an invitation to host the next workshop in early 2009 in Japan. In 2007, the first workshop took place in Cape Town, South Africa, this time in Geneva, Switzerland. Having the next workshop in Asia would follow the GEO rotation principle.

Acknowledgments

This workshop was supported and endorsed by the Meraka Institute, CSIR, Pretoria, South Africa, the GEO Secretariat, Geneva, Switzerland, the European Commission through the SANY and OSIRIS projects (both are Integrated Projects, contract number 0033564 and 0033475, co-funded by the Information Society and Media DG of the European Commission within the RTD activities of the Thematic Priority Information Society Technologies), and the Open Geospatial Consortium (OGC).

References

[1] http://www.sany-ip.eu

[2] http://www.osiris-fp6.eu

[3] Simonis, I. (2007), The Sensor Web: GEOSS's Foundation Layer. In: GEO Secretariat (Ed.), The Full Picture, Tudor Rose, pp. 61-63

[4] http://www.52north.org/

[5] http://www.ogcnetwork.net/node/383

[6] http://www.opengeospatial.org/projects/groups/sensorweb

[7] ftp://ftp.wmo.int/Projects/GEO/TaskSheets/2008-05/

[8] http://www.brockmann-consult.de/CalValPortal/welcome.do

[9] http://www.sensorbay.ca

[10] http://www.intamap.org/