



SANY

an open service architecture
for sensor networks



© 2009 by the SANY Consortium

COPYRIGHT NOTICE

The SANY Consortium grants third parties the right to use and distribute all or parts of this document, provided that the SANY project and the document are properly referenced.

THIS DOCUMENT IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS 'AS IS' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS DOCUMENT, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

ISBN 978-3-00-028571-4

Final design & layout by
Craig Stevens – www.craig-stevens.co.uk

Images by
Cover image © Mangia/Fotolia.com
All other image copyrights are owned by their respective authors.

1 <http://www.sany-ip.eu/partners>

Editors

Martin Klopfer, OGCE; Ingo Simonis, OGCE

Contributing Authors

Thomas Bleier, ARC; Bojan Božić, ARC; Roman Bumerl-Lexa, ARC; Alexandre da Costa, ISEO; Sophie Costes, SDA; Ionut Iosifescu, ETHZ; Olivier Martin, IGN; Steven Frysinger, EIG; Denis Havlik, ARC; Désirée Hilbring, IITB; Patrick Jacques, SPB; Martin Klopfer, OGCE; Siegbert Kunz, IITB; Peter Kutschera, ARC; Mark Lidstone, BMT; Stuart E. Middleton, IT-INNOVATION; Zoe Roberts, BMT; Zoheir Sabeur, IT-INNOVATION; Johannes Schabauer, ARC; Sascha Schlobinski, EIG; Tingting Shu, BMT; Ingo Simonis, OGCE; Bernard Stevenot, SPB; Thomas Usländer, IITB; Kym Watson, IITB; Ken Wittamore, Triskel

The SANY Team

Ludivine Berden, SPB; Thomas Berlinghoff , EIG; Ralf Bommersbach, IITB; Arndt Bonitz, ARC; Michael Brill, EIG; Yves Coene, SPB; Ralf Denzer, EIG; John Dering, MTS; Nicola Dewey, BMT; Pascal Dihé, EIG; Phillipé Dostert, EIG; Markus Eisenla, IITB; Julian Fischer, EIG; Benjamin Friedrich, EIG; Reiner Güttler, EIG; Dr. Markus Hasel, KTT-iMA; Louis Hecht, OGCE; Thorsten Hell, EIG; Thorsten Herter, EIG; Marco Hugentobler, ETHZ; Lorenz Hurni, ETHZ; Jürgen Kost, KTT-iMA; Michael Kunz, KTT-iMA; Regina Kunz, EIG; Wolfgang Kunz, KTT-iMA; Allan McVeigh, BMT; Wenjie Ma, EIG; Christian Marchionini, ISEO; Christophe Meynard, IGN; Jost Nielinger, KTT-iMA; Christophe Noël, SPB; George Percivall, OGCE; Thierry Person, IGN; Bogusz Piliczewski, BMT; Sebastian Puhl, EIG; Rainer Roeckle, KTT-iMA; Matthew Rymell, BMT; Nicola Salvo, IT-INNOVATION; Martin Scholl, EIG; Stéphane Slusarek, SDA; Nils Steinbiß, EIG; Jörg Stumpp, IITB; Christian Thom, IGN; Galina V. Veres, IT-INNOVATION; Paul W. Walland, IT-INNOVATION; Jonathan Williams, MTS; Zlatko Zlatev, IT-INNOVATION

Contents

1	Why SANY?	5	6.8	Data Fusion and Modelling	64
1.1	Is this Book for You?	5	6.9	Implementation Architectures for Fusion Services	72
1.2	The Need for SANY	5	6.10	Implementation Architecture for Decision Support	78
1.3	Using this Book	6			
1.4	Enabling the Sensor Web	8			
2	The SANY Project	9	7	SANY Components	82
2.1	SANY Objectives	9	7.1	Sensor Observation Service	82
2.2	Project Approach	10	7.2	Cascading SOS	84
2.3	Key Results	11	7.3	Sensor Planning Service	86
3	SWE – A Global View	12	7.4	Web Notification Service	87
3.1	SWE Initiatives	15	7.5	Web Services Notification	88
3.2	SWE Standardisation	17	7.6	Sensor Alert Service	89
3.3	SWE in A Global Context	17	7.7	Catalogue Service	91
4	SANY's Use Cases and Pilots	19	7.8	Map and Diagram Service	95
4.1	Air Quality Management	19	7.9	Web Processing Service	97
4.2	Decision Support Tools for Marine Risk Management	20	7.10	Time Series Toolbox	99
4.3	Geo-hazards in Dense Urban Areas	22	7.11	SensorSA Data Acquisition System	101
5	SANY Value Proposition	25	7.12	Generic Time Viewer	105
5.1	Why should you consider using SANY results?	26	7.13	IGN GeoCubes	107
5.2	How to use the results of SANY?	27	7.14	SolData Microns	111
6	The Sensor Service Architecture	28	7.15	Triskel Marine Ltd. DataBuoy	113
6.1	Design Principles	31	8	SANY Applications	116
6.2	Reference Model for Open Distributed Processing	33	8.1	Air Quality Management	116
6.3	Sensor Models	34	8.2	Decision Support Tools for Marine Risk Management	129
6.4	Functional Domains	38	8.3	Geo-hazards in Urban Areas	138
6.5	Security Framework	40	8.4	Conclusions	145
6.6	The Standards State-of-the-Art	46			
6.7	Elements of the SensorSA	52	Annex 1: Abbreviations and Acronyms		147
			Annex 2: Glossary		150
			Annex 3: Bibliography		157



Why SANY?

1

IS THIS BOOK FOR YOU?

1.1

- Do you deal with civil contingencies and environmental risk management?
- Are you monitoring our fragile environment, using various sensors?
- Does your data and information content deserve exposure to broader markets?
- May your domain expertise be needed in a time of emergency?
- Are you implementing parts of a service chain related to any of the above?

If your answer to any of the questions is yes, then perhaps you are keen to understand where current trends in technology and society are taking us and how these trends impact our life by helping to build an increasing awareness of environmental issues. If this is the case, then yes: this book is for you!

And should this introduction sound familiar, than you might very well be correct: this book summarises the approaches and results of the SANY project by following the example of the ORCHESTRA project, whose work on an open architecture for risk management has provided the foundation for SANY. The acronym SANY stands for Sensors Anywhere and embraces trends and approaches identified by ORCHESTRA, many of which have by now developed into reality. As a major Integrated Project in the Sixth Framework Programme of the European Commission, SANY extends the work of ORCHESTRA into the domain of sensor networks and standards based sensor web enablement.

THE NEED FOR SANY

1.2

Our very own daily life and routines are constantly influenced by environmental aspects, and, whether consciously or unconsciously, we react to these impacts and adjust our own activities accordingly. Whilst this sounds a lot like common sense, rather than the rational for another project to be funded by the European Commission, there's a deeper layer of relevance to this: our own understanding as individuals of our environment, as well as the common understanding as a society of potential environmental threats to our way of life has improved tremendously over the past decades. Be it solar radiation, ozone levels, fine particulate matter exposure, bathing water quality or more subtle topics, such

as subsidence of buildings due to infrastructure development: most, if not all, of these environmental impacts on our lives are known on a broader scale. It is this growing understanding, which eventually helps to protect the environment and promote conditions, which are beneficial to the current generation and as well as those to come.

So, what about sensors anywhere?

An observation leads to information; information leads to knowledge, to understanding and ultimately understanding may even lead to the wisdom to act accordingly.

It is this very chain that leads from abstract ozone measurements to a common wisdom to ban CFCs from widespread household usage. This is the very point where SANY contributes: making observations from sensors available in a more readily, widespread and interoperable fashion will help to improve our understanding of environmental processes and impacts on our life. It will also support the development of fusion, interpretation and visualisation tools that provide the base for well informed, improved decision making.

SANY has identified and addressed the major technological challenges and barriers for efficient information handling between stakeholders. This includes a number of different scenarios, where sensor data is the starting point for decision making processes in the domains of air quality management, geo-hazard mitigation and coastal water quality control.

1.3 USING THIS BOOK

This book is broadly split into two mostly self contained parts and we suggest you start with the part you feel most comfortable with:

1.3.1 The Business Perspective

The initial chapters of this book provide a quick summary of the basic approach, key results and general benefits of the SANY project.

Reading this section will provide you with:

- an understanding what the SANY project is about
- business reasons for adopting an open standards based architecture approach
- guidance on next steps to improve your own future projects
- pilot examples on how SANY results contributed to real life scenarios

The Technical Perspective

1.3.2

The second part of this book discusses in more detail the approach and results of SANY. It puts a strong emphasize on introducing the concept of the SANY Sensor Service Architecture as well as specific sensor services, which have been developed and/or deployed in the SANY pilot implementations. This part of the book is targeted towards technically minded readers who seek entry-points to understand and develop their own standards based sensor networks as part of a larger interoperable sensor web.

This section will give you:

- the information needed to build your own sensor service network
- the information on how SANY Pilots implemented sensor services
- information on all major services, software components and related standards
- access to software components developed and used by SANY

SANY has worked very closely with a number of Working Groups of the Open Geospatial Consortium (OGC), adopting existing standards for pilot implementations and feeding back requirements for standard extensions and improvements into the global standardisation process.

The SANY pilots specifically highlight the benefits of being able to task and query sensors through interoperable networking, rather than having to rely largely on proprietary arrangements. Being able to use sensors as well as services in an interchangeable and interoperable fashion boasts a whole range of new opportunities for information collection, research and subsequent business development.

This section of the book is complemented by a set of tutorials and open source licensed software components, which are available for download online at the SANY website.

1.4 ENABLING THE SENSOR WEB

As ORCHESTRA already highlighted, access to relevant information is one, if not *the* most relevant improvement in the highly complex network of environmental risk management. Sensor data is the most direct link we can have to monitor and analyse changes in our environment and correlate the results with likely impacts on society.

More and more sensor data sources become available, but only when they describe and communicate their capabilities and observations through interoperable standardised interfaces will our understanding of our environment and its potential impact on our life improve further.

Likewise, following the standardised interface approach will help to deploy the full potential of a sensor network and through its versatility to adopt to future tasks help to protect the initial investments in its deployment.

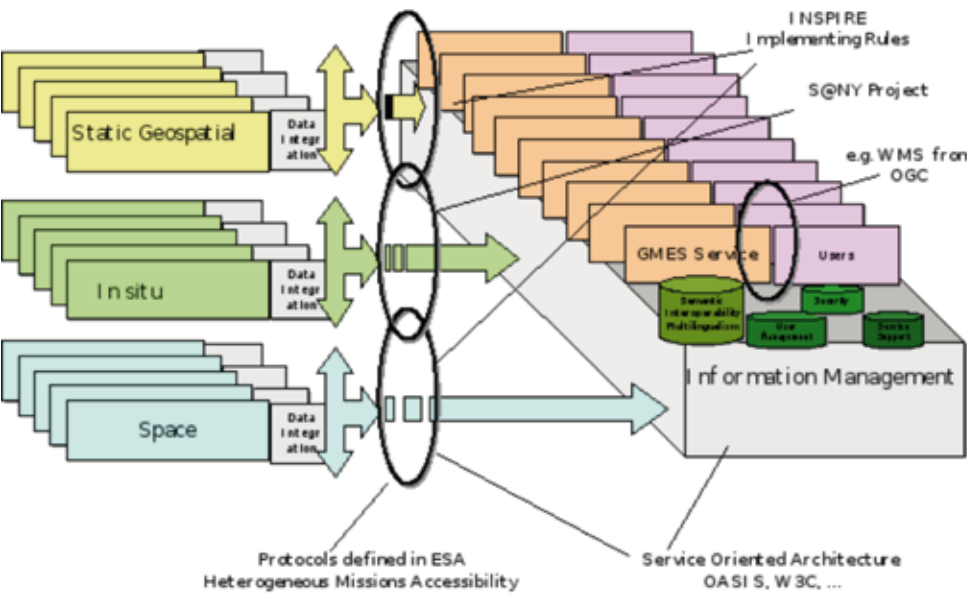
INSPIRE, GMES and GEOSS are well known examples of activities that aim to improve decision making and governance on a multi-national scale based on information that relies on a whole range of sensor data at all scales, from in-situ as e.g. for ozone concentrations in ambient air, to earth observation data to determine e.g. land cover classes. By adopting and promoting the use of standards and feeding back identified requirements to the respective standards organisations, SANY has set a best practise example, whose adoption will further boost the success story of open and standards based service architectures. The standards based approach to sensor web enablement helps organisations on all levels of involvement to flexibly adapt their networks and services to potential new requirements.

The SANY Project

SANY OBJECTIVES

2.1

SANY aims to improve the interoperability of in-situ sensors and sensor networks, allowing quick and cost-efficient reuse of data and services from currently incompatible sources in future environmental risk management applications. Whilst INSPIRE addresses largely access to static geospatial data and the Heterogeneous Mission Accessibility (HMA) initiative of the European Space Agency addresses earth observation data, access to and interoperability between in-situ sensors has not yet been specifically addressed. The graphic below outlines the positioning of the SANY project versus the core areas of INSPIRE and GMES and highlights how the work of SANY helps to pave the way from data to information:



Based on: GMES Reflection paper on Data integration and information management capacity, DG-INFOS, Draft 6, July 2005; this diagram is slightly modified in order to illustrate the positioning of the SANY project

The main objectives of SANY are:

- to specify a generic open architecture for fixed and moving sensors and sensor networks capable of seamless ‘plug and measure’ and sharing of resources in virtual networks;
- to develop and validate reusable data fusion and decision support service building blocks and a reference implementation of the architecture;
- to closely work with end users and international organisations in order to assure that the outcome of SANY contributes to future standards;
- and to validate the project results, through development of three innovative risk management applications covering the areas of air pollution, marine risks and geo hazards.

2.2 PROJECT APPROACH

Commercially a project like SANY only makes sense when its results address the targeted users’ needs – whilst a significant amount of research work was undertaken, which may not directly impact today’s IT solutions, there’s a strong interest and commitment of the consortium partners to engage with potential system users and stakeholders at an early stage and involve them in the *design – build – validate* cycles which have also been deployed in the SANY Pilot implementations.

Due to its strong links to ORCHESTRA, the project adopted the *OGC approved Best Practise Reference Model of the ORCHESTRA Architecture* (RM-OA) as a starting point. The RM-OA helps to identify user requirements and translates them into generalised specifications. Based on these requirements, SANY’s core work is the development of the Sensor Service Architecture (SensorSA), reference services, data fusion and modelling services, and generic building blocks for decision support applications.

To ensure that developments meet exploitation requirements, the project followed an iterative approach of 3 cycles of the following steps, in which the results of each completed cycle were used to further refine the requirements for the following phase:

- identification of user requirements and available complementary project activities,
- development of system and architecture specifications,
- implementation of pilot systems,
- validation by end-users.

KEY RESULTS

2.3

If you wish to follow the example of SANY and want to establish your own nodes or branches of the sensor web, you may find one or more of the following public documents of interest, which summarise the experiences and results of the SANY approach and may thus help to define a realistic project plan and overcome initial hurdles:

1. Sensor Service Architecture (SensorSA) specification
2. Prototype implementation of the SensorSA services
3. a framework for integration of fusion- and modelling- engines into SensorSA networks
4. a security framework for access control & policy enforcement
5. a web based platform for decision support applications based on ESA SSE
6. three prototype applications illustrating the use of SANY in air quality, marine risks and geo hazards domains.
7. a collection of educational material for decision makers and technicians interested in developing their own SensorSA compliant networks has been created and is available together with required open source software components online at the SANY website.

Beyond the collection of reports and materials, which has been created in SANY, the most valuable outcome is probably the experience of the joint engagement in using and extending standards for defined pilot use cases.

Discussing requirements not only amongst project partners, but in the community of likeminded experts with an interest in sensor web enablement often helps to develop new ideas to approaches and solutions. It’s the networking aspect of the stakeholders which is probably as important as the networking of sensors and services.

Whilst all efforts have been made to summarise the experiences and results in this book, the complementary tutorials and the public deliverables, you shouldn’t hesitate to contact the SANY Consortium for further information or support when needed.

3 SWE – A Global View

Sensors are everywhere. You find them in your house, a supermarket, in streets, bridges, rivers, on and in oceans, air and space. They measure various phenomena, like the temperature in your refrigerator, the pressure in pipelines, if someone approaches (to open a door, light the yard etc.), the height of waves, water quality, building stress, and many, many more. We use them for various purposes: surveillance, monitoring, prediction, controlling and often for our convenience (think of a GPS in car navigation).

A tremendous amount of information is generated each second but we are far from tapping its full potential.

Why? There are several reasons:

- First of all, the sensors and networks of sensors are usually disconnected, meaning that they are not connected to a globally accessible information network.
- Second, even if they are connected, we usually do not know how to search for those sensors that are of interest to us.
- Finally, even if we have found new sensors, we cannot easily make sense of the data provided by them, due to their proprietary data interfaces and encodings. Luckily, there is a solution to these challenges: it is called the Sensor Web.

The Sensor Web started as a conceptual design study several years ago. Today, though far from being complete, it is instantiated. Hundreds of sensors and other components already contribute to the Sensor Web and the number is continuously growing. So what constitutes the Sensor Web?

In the Sensor Web:

- Sensors and sensor networks are connected and accessible via the World Wide Web.
- Access to sensor information and observations will be achieved through standardized Web service interfaces.
- Sensors are self-describing to both humans and software alike, using standard (non-proprietary) encodings.
- Thus, these sensors and ultimately their data will be discoverable. Much like search engines are capable of finding content in web pages across the

globe, the Sensor Web provides components to search for specific sensors and sensor data – of the past, present and future.

- Through standardized Web service interfaces, sensors, simulations, and models will be capable of being configured and tasked dynamically.
- Software will be able to geolocate and process observations from newly discovered sensors without a priori knowledge of the sensor system that generated the observations.
- New and higher-level information will be generated on-the-fly based upon the vast source of sensor data now available.
- All this information will be distributed and alerts be raised when events of interest are detected, enabling the initiation of responsive action, even automatically.
- Sensors will be able to act on their own (i.e., autonomous), even in concert, based upon the rich offer of information about their environment.

As all components of the Sensor Web (such as sensors, access interfaces, data stores etc.) are operated and maintained by different organizations, it is a set of common agreements that bootstraps the Sensor Web. Standardisation organizations coordinate the process of finding common ground and mutual agreements among experts and sensor operators from various domains. The goal is to develop a framework of standards generic enough to support a wide field of applications while remaining specific enough to ensure interoperability among all participating components.

The Sensor Web builds on the World Wide Web and uses a wide variety of standards recommended by the W3C, such as XML, XML Schema or SOAP for data encodings and interface specifications. Using the Web as its foundation layer, the Sensor Web makes use of Web technologies and supports the integration of communication infrastructures taking place on lower levels of the communication stack, often using standards developed by the Institute of Electrical and Electronics Engineers (IEEE) or the Internet Engineering Task Force (IETF).

Thus, the Sensor Web is a middleware layer that enables access to sensors and sensor data using Web technologies. The Sensor Web standards themselves are mostly developed by the OGC.

Since 2001, the standards developed by the OGC working group ‘Sensor Web Enablement’ (SWE) have matured and have now reached a stage where the first version of the Sensor Web can be implemented. So what is SWE exactly?

It is:

- A technology to enable the realization of the Sensor Web, much like TCP/IP, HTML and HTTP enabled the World Wide Web.
- A suite of open, consensus-based standards defining encodings and Web service interfaces required in the Sensor Web.
- A Service Oriented Architecture (SOA) approach, so it integrates with mainstream IT approaches.

SWE supports the integration of virtually any sensor technology into the Sensor Web. It can be used with and applied in restricted sensor networks but also in medium and large scale or global networks (like INSPIRE, GMES, or GEOSS, see chapter 3.4).

Because of its service oriented approach, it enables distributed architecture development and deployment as well as on-the-fly connectivity between resources. Care has been taken that SWE facilitates incremental migration of existing proprietary sensor networks into the Sensor Web. SWE makes use of standard models and also semantic concepts, which ultimately enables interoperability in the Sensor Web. In addition, the technology supports up-to-date IT mechanisms to ensure security and scalability of the infrastructure.

The specifications that comprise the Sensor Web Enablement suite of standards developed by the OGC are presented in the following:

- SWE Common – specifies a data model and encoding to define and package sensor related data in a self-describing and semantically enabled way. It is used by several other SWE standards.
- Sensor Model Language (SensorML) – defines a data model and encoding to describe processes and processing components associated with the measurement and post-measurement transformation of sensor observations.
- Observations and Measurements (O&M) – defines a data model and schema for encoding measurements and observations.
- Sensor Observation Service (SOS) – defines a service model and interface encoding for the provision of sensor measurements and observations, from simple sensors to complex sensor systems, both physical and virtual.
- Sensor Planning Service (SPS) – defines a service model and interface encoding for the execution of sensor tasking and parameterization requests.

- It is used to manage sensors and sensor networks and to influence the measurement process according to specific needs and requirements.
- Sensor Alert Service (SAS) – defines a service model and interface encoding that enables subscription for and notification of situations of interest based upon continuous evaluation of incoming sensor observation streams.
 - Web Notification Service (WNS) – defines a service model and interface encoding for distributing incoming information to registered users via various communication protocols. It is often used for supporting asynchronous communication and routing urgent messages to whole groups of users according to their communication preferences.

SWE INITIATIVES

3.1

Many projects and initiatives apply SWE to integrate their sensors and sensor networks with the Sensor Web. They have helped to mature SWE technology and are the means to continuously improve the existing specifications.

EU R&D Projects and Initiatives

3.1.1

Several projects funded by the European Commission and other European organizations further SWE and apply the technology in their real world use cases. Projects like SANY and OSIRIS show that SWE can be applied in various risk monitoring and risk management scenarios in multiple societal benefit areas that are also relevant for GEOSS (see chapter 3.4).

The European Space Agency (ESA) initiated in 2005 the Heterogeneous Mission Accessibility (HMA) project. ESA and various partner organizations in Europe, are collaborating on the objective to harmonise access to heterogeneous earth observation missions, including national missions and ESA Sentinel missions. HMA involves a number of OGC standards, including the Sensor Planning Service, which supports the feasibility analysis requirements of Spot Image optical satellite missions.

Many more initiatives exist that apply SWE in various domains, ranging from defence and intelligence over tsunami early warning to home automation.

National Initiatives

3.1.2

There are several national and international initiatives under way implementing Sensor Web components in order to address various challenges in an efficient way. One example is the nationwide Water Resource Observation Network

(WRON), an Australian flagship project striving to improve Australia's water information leading to the improved management of water resources and to establish the technological platform for integrated water information systems across Australia. Among the most advanced components of WRON, we find the Tasmanian Hydrological Sensor Web. Here, the Commonwealth Scientific and Industrial Research Organisation (CSIRO), Australia's national science agency are building a 'hydrological Sensor Web' covering the South Esk catchment in North East Tasmania. The Tasmanian Hydrological Sensor Web will integrate rainfall, climate and stream flow data collected by in-situ sensors with numerical models that produce daily quantitative precipitation forecasts, rainfall-runoff estimates and stream flow predictions.

Another example is the Advanced Fire Information System (AFIS) in South Africa. Here, Sensor Web technologies are used to detect and alert about devastating wild fires. Operated by the Council for Scientific and Industrial Research, AFIS gets continuously enhanced to serve as a fire information system for sub-Saharan Africa.

The US National Aeronautics and Space Administration (NASA) has adopted the vision of sensor webs as a strategic goal and has thus funded a variety of projects to advance Sensor Web technology for satellites. Central to many of these efforts has been the collaboration between the NASA Jet Propulsion Lab and the NASA Goddard Space Flight Center using the Earth Observing 1 (EO-1) and assorted other satellites to create pathfinder Sensor Web applications, which have evolved from prototype to operational systems.

Sensor Web concepts are further explored by the private industry. As an example, Northrop Grumman Corporation (NGC) has been using the SWE standards in a major internal research and development (IRAD) project called Persistent Universal Layered Sensor Exploitation Network (PULSENet™). This real-world testbed's objective is to prototype a global Sensor Web.

3.1.3 OGC Testbeds

Regular testbed activities conducted by the OGC with participation of organizations and individuals from across the globe led to the current status of SWE. From the first Open Web Services (OWS) testbed 1.1 to the recent OWS-6, SWE has always had its place in the various successful capability demonstrations.

3.1.4 OGC Interoperability Experiments

In addition to testbeds, the OGC performed (and performs) several interoperability experiments (IE) to further certain standards (like the Sensor Alert Service) or to investigate the applicability of SWE standards for a given

application domain. The most recent IE (Oceans IE) applied SWE in the context of oceanography and ocean communities. These efforts will continue (Oceans IE phase II).

SWE STANDARDISATION

3.2

Sensor Web Enablement is a standardisation effort driven by the Open Geospatial Consortium. Through its liaison with ISO TC211, the OGC is closely involved in the development of often legally binding services published by ISO. To date, OGC has successfully established its Observation and Measurement specification as a new work item in ISO. In the future, other SWE standards will follow.

The development of SWE by the OGC sometimes seems to collide with efforts from other standards organizations, such as the IEEE 1451 family of standards, which also deals with sensor networks and their uniform connection to a greater network. The OGC SWE group has always collaborated with these efforts, exchanging knowledge and performing combined testing activities. This helped to clarify the role of SWE and to improve the standards.

Many IT technologies are being subject to standardization efforts, especially when service oriented architectures are concerned. When mainstream IT standards are concerned, the IETF, W3C and OASIS are important standardisation organizations. The OGC SWE group is working with these organizations on different aspects. OGC's focus is to geospatially enable mainstream IT, because location plays a vital part in most of our daily activities. Thus the members of OGC pay close attention to new developments and adopt and apply approved IT standards where applicable.

SWE IN A GLOBAL CONTEXT

3.3

GMES (Global Monitoring for Environment and Security) is a European initiative for the implementation of information services dealing with environment and security. It is based on observation data received from Earth Observation satellites and ground based information. These data will be coordinated, analysed and prepared for end-users in order to understand the short, medium and long-term evolution of the environment and to help European citizens to improve their quality of life. Built up gradually, GMES is one of Europe main contribution to an even larger initiative: GEOSS, the Global Earth Observation System of Systems.

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. Implemented as an emerging public infrastructure to interconnect a diverse and growing array of instruments and systems for monitoring and forecasting changes in the global environment, GEOSS addresses a number of societal benefit areas, as there are disasters, health, energy, climate, water, weather, ecosystems, agriculture, and biodiversity. The integration and often timely delivery of earth observation data is a key to all of them.

The Sensor Web presents a paradigm in which the Internet is evolving into an active sensing macro instrument – an instrument capable of bringing sensory data from across the globe to the finger tips of every individual. Thus it is no wonder that SWE standards play a major role in the emerging GEOSS infrastructure.

SANY’s Use Cases and Pilots

4

AIR QUALITY MANAGEMENT

4.1

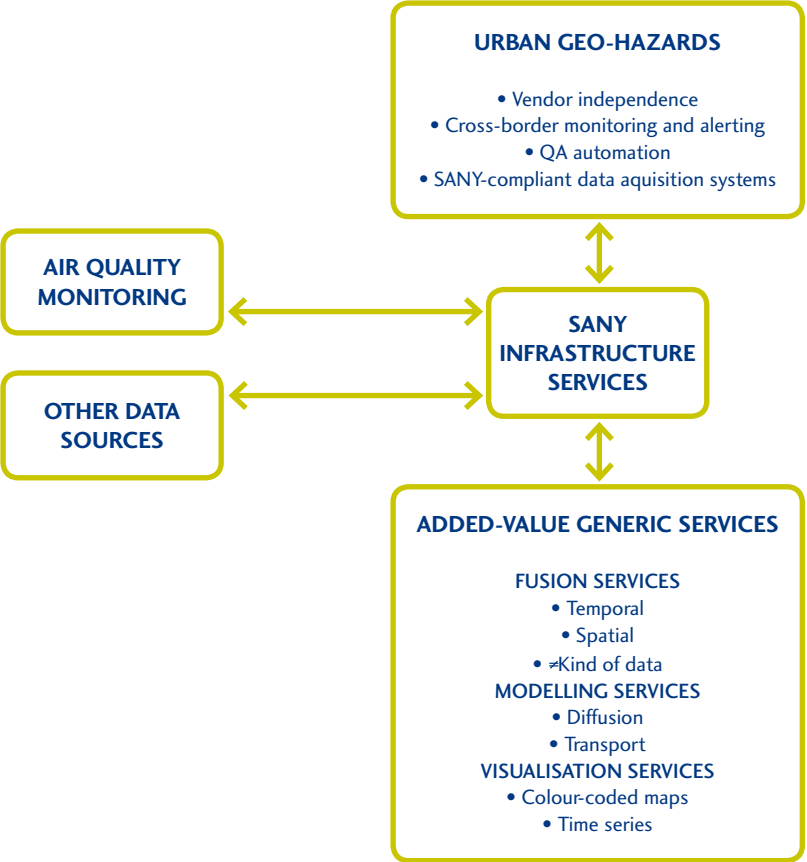
Air quality is one of the most important indicators for the sustainable development. The air quality monitoring is therefore required and regulated by the law in all European states. In addition, to existing reporting obligations the EU-wide initiatives such as INSPIRE and CAFE are gradually introducing the need for the pan-European interoperability and real time exchange of data.

The SANY ‘Air Quality’ pilot is used to validated the usability of the SensorSA based air quality management networks for three main groups of users: network operators, national environmental agencies, and for the European Environmental agency.

The SANY Air Quality Management Pilot focuses on the following topics:

- Providing uniform access to data from air quality monitoring systems of France, Belgium and Austria. The Air Quality Pilot also showcases the feasibility of serving the INSPIRE-relevant meta information over the standardized OGC Sensor Observation Service interface
- Aiding the domain experts in performing the routine Quality Assurance of the data. This is achieved by mean of the state space fusion service. This service continuously monitors all available air quality (immission) observations and publishes the now-casts and confidence intervals at 17 measurement locations using the data model similar to the original immission data model. The combination of the data from both servers, presented side-by side provides a very effective help in finding suspicious measurements.
- Identifying the impact of the known pollution sources to actually measured immission, and providing an indication for the relative importance of the unknown (unaccounted for) sources of pollution at the selected positions. This is achieved by comparing the immission measurements with the prediction based on real-time emissions from major industrial plants in the Linz area.
- Illustrating the feasibility of the automatic report generation. This use case is limited to automatic generation of the data required for reporting in the CAFE

The illustration below shows a summary of the main components of the air quality applications.



4.2 DECISION SUPPORT TOOLS FOR MARINE RISK MANAGEMENT

Marine risks arise from a number of sources including natural events, anthropogenic causes and a combination of both. In almost all cases, marine risks have an economic, human and environmental impact.

In SANY, short term microbial contamination of both Bathing Waters and Shellfish Waters has been targeted. These designated water areas are subject to extensive regulatory standards, established via EU Directives.

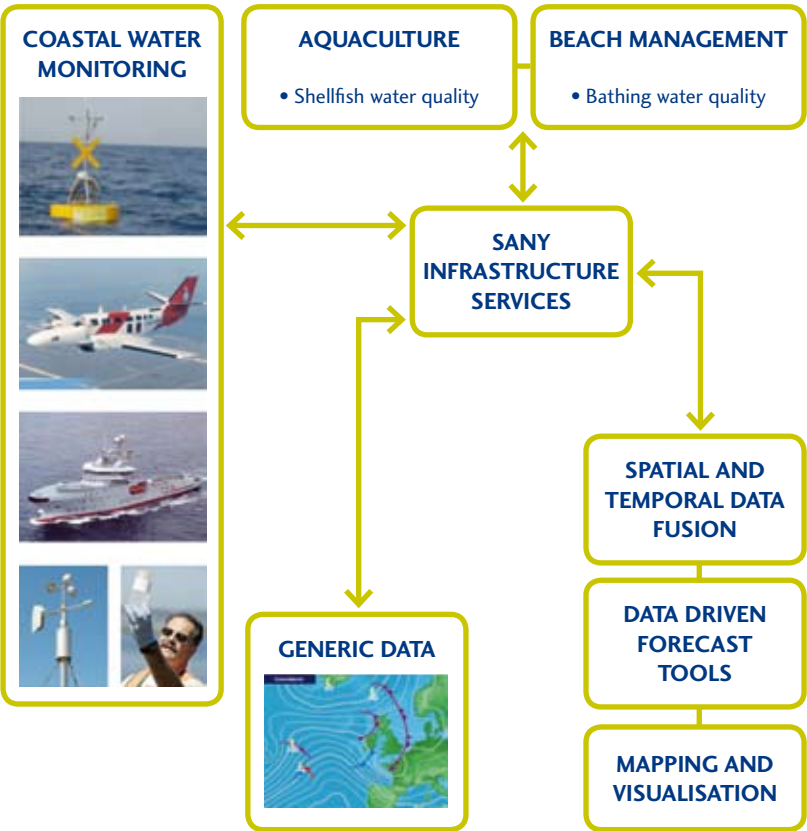
In the case of Bathing Waters, failure to meet regulatory standards can have significant impacts on public health and tourism revenue. Similarly,

short term microbial pollution events in Shellfish Waters can have serious consequences for consumer health and cause a reduction in revenue for the local aquaculture industry.

Presently, microbial levels in the selected water areas are assessed using laboratory testing. These tests often have a turnaround time of more than 24 hours and, as such, can only determine whether a contamination event has occurred. Improvement in the ability to forecast the risk of short term microbial pollution in designated waters could reduce both the human and economic impact of such events. The SANY marine risk applications use a number of services developed within the SANY project to access data from sensor networks and assess the likelihood of a contamination event occurring. The use of SANY Sensor Service Architecture enables:

- Access to third-party sensor networks and phenomenological models, to create cost-effective access to measured or modelled data streams and equally to allow operators of such networks/models to valorise their investments;
- The use of web-based services to provide high-value data processing (eg for spatial fusion, temporal fusion and modelling) that will enable users to get enhanced information about parameters of interest;
- Rapid deployment of additional in-situ sensors on both fixed and mobile platforms. These will acquire data on key water quality parameters, to fill gaps in spatial and temporal data coverage, and thereby permit improved quality of risk now-casting and forecasting;
- Provision of alerts and alarm systems to raise the awareness on a possible hazard and support preventative measures;
- Remote configuration of smart sensors and, if possible, adaptive tuning of stochastic models to allow ‘on the fly’ enhancement of risk forecasting through incorporation of recent data within the forecasting algorithm.

The illustration below shows a summary of the main components of the marine risk applications.



Section 9.2 provides further details on the Bathing Water and Shellfish Water risk applications that have been implemented as SANY Pilots.

4.3 GEO-HAZARDS IN DENSE URBAN AREAS

Geo-hazards may be caused by human activities or natural events. Whether those hazards are induced by human activity or natural hazards, they have an economic, human and environmental impact, which cannot be neglected. As an example, landslides are among the most widespread hazards on Earth causing billions of dollars in damage and thousands of deaths and injuries each year around the world, and Europe has the second highest incidence of landslide

casualties of any other continent. As well, recent accidents in European cities induced by construction works raised the awareness for a better control of monitoring data, and enhanced services for decision support.

In SANY, the geo-hazards pilot focuses on hazards related to construction works in dense urban areas. Indeed, with the expansion of urban areas and the densification of population and transport networks of those areas, construction and rehabilitation works on structures have become more frequent, thus the population is exposed to higher risks. There is therefore a critical need for a better management of geotechnical risk in such a context.

Moreover, monitoring systems and sensor management software installed on a construction site are usually proprietary, and vary from one provider to the other, thus multiplying data sources and information. With respect to those limitations, the Geo-hazard application intends to provide an easy and fast access to sensor data, independently from the sources, and the possibility to merge that information through fusion and modelling services, in order to offer synthetic and comprehensive information to the end-user.

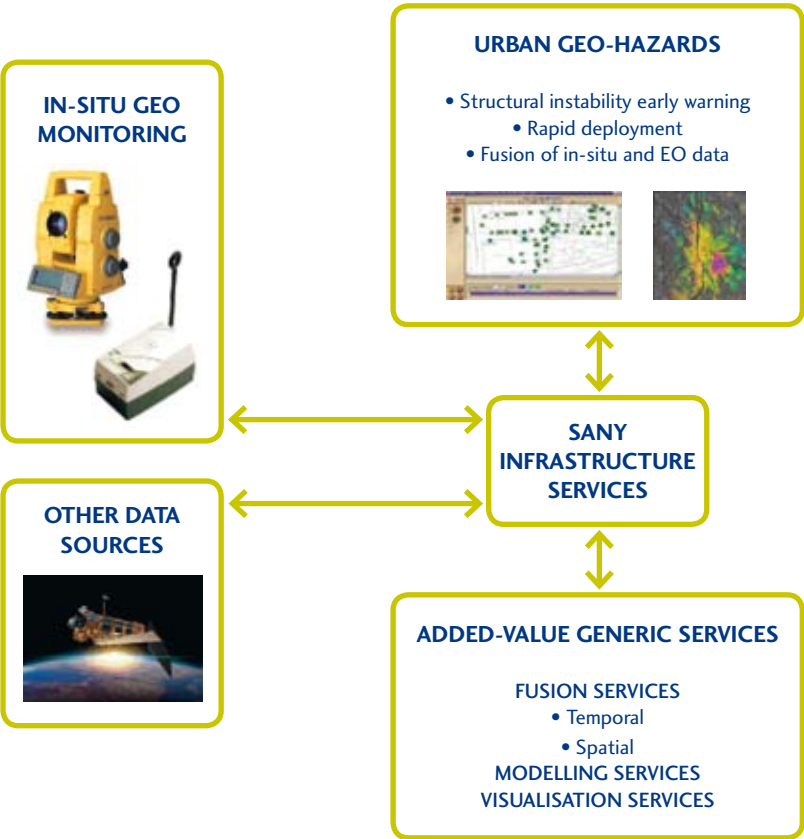
Although the SANY geo-hazard application focuses on the risk management in urban areas due to construction works, most of the services used and implemented for this Pilot may be transposable and used in other contexts (landslides, structural health monitoring, etc).

The use of SANY Sensor Service Architecture enables:

- A common and interoperable access to third party in-situ, EO data, and wireless smart sensors data for a more comprehensive and global information;
- A compliance of information between different systems using well-defined resources identifiers, as well as a standard description of sensors and sensor systems;
- The provision of alerts when alarms conditions are met, and a customised notification of such alerts by the user for a better awareness on a possible hazard and support preventative measures;
- The remote configuration and management of wireless smart sensor networks;
- The Fusion of distributed in-situ measurements of geophysical parameters with other relevant data (e.g. EO data, topographic data, ...) in order to generate more accurate information;
- The provision of an early risk awareness information using predictive services (temporal fusion and the use of geotechnical models) to predict alarms and ensure a faster response to a potential risk;

- The possibility to have additional information where no sensor measurement is available through spatial fusion services or through the rapid deployment of additional in-situ sensors;
- The use of a Services platform onto which the SANY services are grafted. The services are chained to one another, using a workflow engine that triggers the services and passes the information in a standardised way from one service element to the other, in order to create new applications that will be used for monitoring and forecasting environmental geophysical phenomena.

The illustration below shows a summary of the main components of the geo-hazard application:



Section 8.3 provides further details on the Geo-hazard application that have been implemented as SANY Pilot.

SANY Value Proposition

5

The value proposition of SANY is directly related to the current issues which need to be resolved in order to allow for a long term sustainability of FP6 and FP7 RTD works in the domain of Environment, Health, Security and Risk Management, i.e.:

- well organized, seamless access to information
- security of access to data and safety of data repositories
- reliability of access to data
- confidence of information contents and service performance
- economic model enabling a win-win approach between stakeholders
- ownership and property rights on data and knowledge

Obviously, SANY must look far beyond the Research and Technology focus and take into account the other dimensions of the challenge of dealing with information exchange:

The exchange of information, whatever the nature and purpose of use of the information, is an economic process; therefore, a set of mechanisms (legislative, financial, organizational, psychological) are needed to facilitate and rule the information exchange.

The demand by stakeholders (Authorities, Organizations, Enterprises, Scientists, Citizens) for reliable, cost-effective, ready-to-use Information related to Environment, Health, Climate, and associated risks can only be achieved by the creation of a socio-economic context which triggers the creation of an open 'Marketplace' of such Information and Services.

SANY provides 'building blocks', which will contribute to establish socio-economic and organizational mechanisms:

- Facilitating the creation of added-value services
- Motivating 'actors' (SME's, Research Institutions, Public Organizations, large industrial Companies) to develop and market thematic and generic web-based services ?
- Motivating 'Data owners' to market their Data.

5.1 WHY SHOULD YOU CONSIDER USING SANY RESULTS?

There are two main reasons why stakeholders involved in Security and Environmental Risk Management have to take the SANY results into account in the evolution of their Information Systems.

- 1. SANY has focused the development work on improving the access to sensor measurements in a coordinated approach with complementary development works dedicated to Earth Observation sensors, Fusion and Modelling, Orchestration of services, Visualization, etc., thus opening the offer of technologies needed to satisfy the communications and interoperability requirements of Risk Management systems. This was made possible because most SANY partners are also involved in key research projects (FP6-ORCHESTRA, ESA-SSE, ESA-HMA, FP7-GIGAS, etc) and contribute actively to the emergence of the interoperability standards (OGC, OASIS, INSPIRE).
- 2. SANY partners are willing to continue the development of these key technologies in a sustainable way and with a coherent vision of providing a response to the challenges, which decision-makers, data providers, users are facing when dealing with natural and man-caused disasters, Environment, Public health, Security etc. Most of these stakeholders deserve robust, flexible, scalable Information and Communications Technologies to overcome the current technical barriers of legacy systems as well as organizational and legal barriers to sharing Information.

The results of SANY are best fit for situations where the stakeholders require synthetic information resulting from the combination of multiple heterogeneous sources of data.

The following use cases illustrate possible domains of application of SANY:

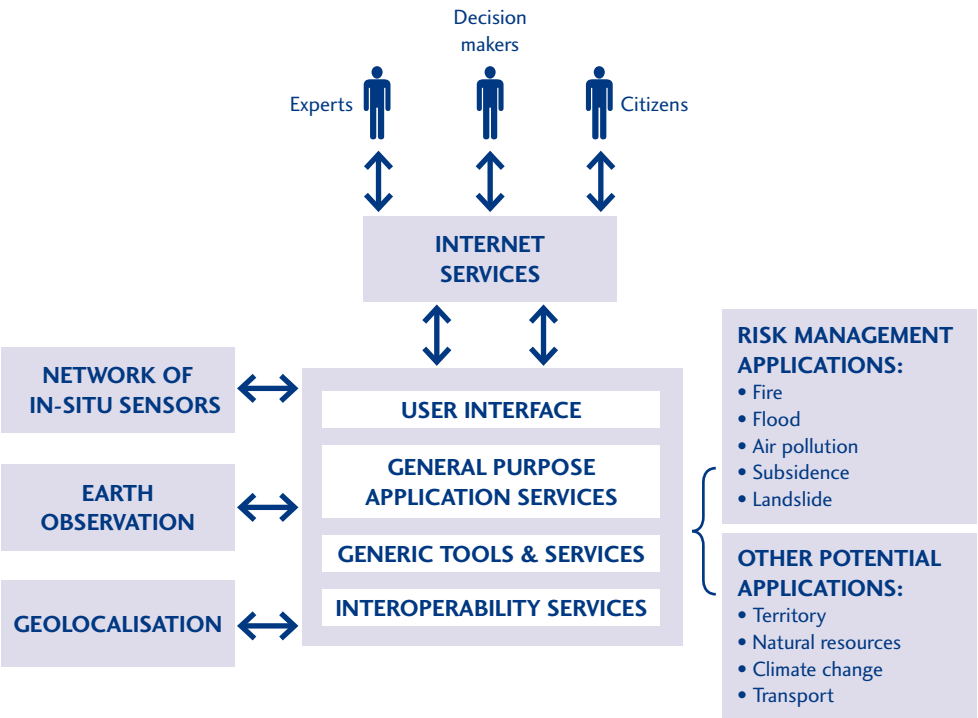
- Public data (water level in rivers, land use, digital maps) should be made seamlessly available to users who need them;
- Citizens suffering from respiratory weakness should be notified, on demand, in case of air pollution surge;
- Emergency services, in case of industrial accident, need up-to-date information about site context (cartography, neighbourhood), products (nature, toxic effects), atmospheric parameters, soil parameters
- Scientists need in-situ measurements combined with clinical data base information in order to perform epidemiological studies.

HOW TO USE THE RESULTS OF SANY?

5.2

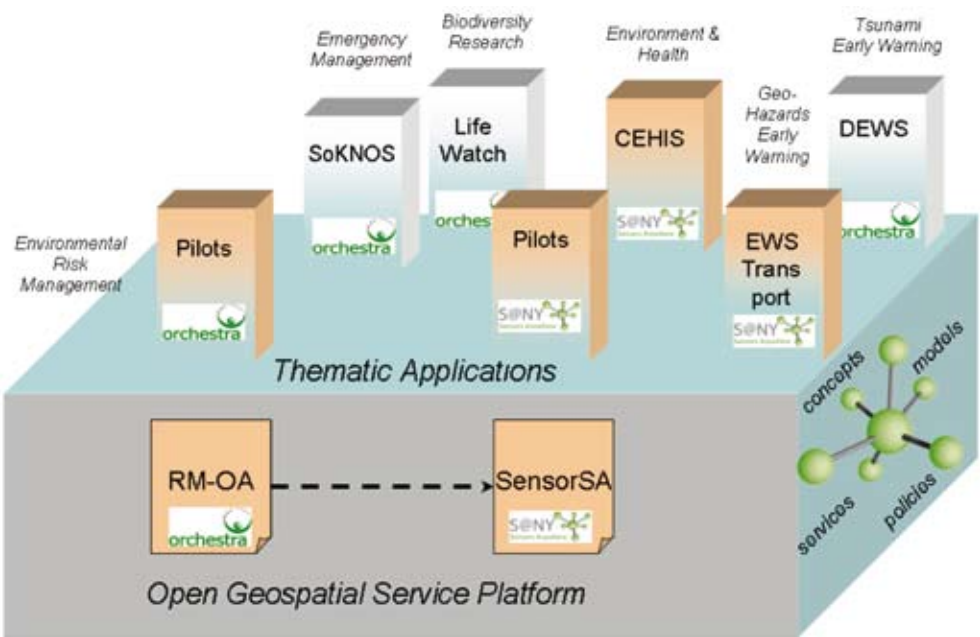
Based on the Service Oriented Architecture and the ORCHESTRA Reference Model, SANY offers the flexibility to tailor the implementation of SANY to the context of the user, ranging from the invocation of a web service up to the creation of an open platform enabling the trade of information and added-value services.

The figure below illustrates the central role of SANY in versatile multi-services platform.



6 The Sensor Service Architecture

The Sensor Service Architecture (SensorSA) is the fundamental architectural framework of the SANY project for the design of sensor-based environmental applications and their supporting service infrastructure. The SensorSA belongs to the family of service-oriented architectures (SOA) with additional support of event processing and a particular focus on the access, management and processing of information provided by sensors and sensor networks. As such, it contains sensor-specific services. However, in order to provide a higher-level interface to environmental risk management applications that is functionally and semantically richer, it abstracts from the peculiarities of sensors and encompasses generic information processing functionality. Thus, there is a gradual transition to the functionality of a generic service infrastructure.



The foundation for the SensorSA is the SOA approach specified by the European Integrated Project ORCHESTRA¹ (Open Architecture and Spatial Data Infrastructure for Risk Management) in its Reference Model for the ORCHESTRA Architecture (RM-OA) (Usländer (ed.), 2007) as well as the OGC Sensor Web Enablement architecture (Simonis (ed.), 2008).

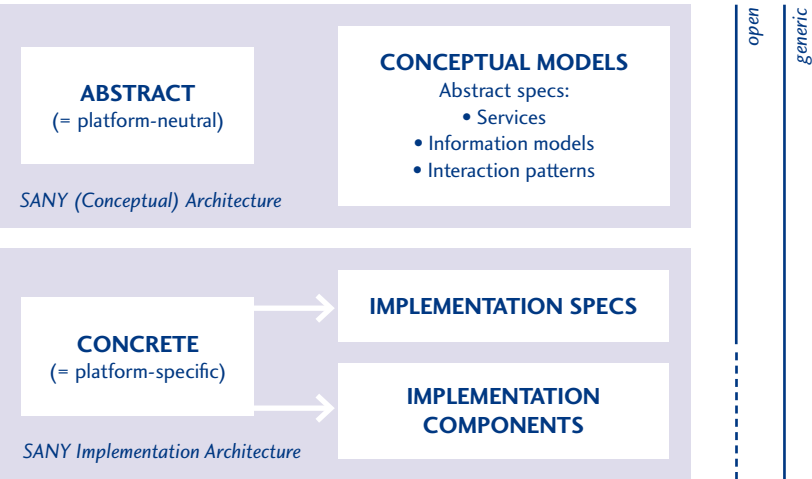
Based upon these architectural frameworks, the SensorSA enables the set-up of open geospatial service platforms for a multitude of thematic applications in different domains. As illustrated by the white boxes in the figure on the previous page, the RM-OA has already been applied and extended in several environmental risk and emergency management applications beyond its use in the ORCHESTRA pilot applications. It serves as foundation for the LifeWatch² reference model supporting the development of e-Science and technology infrastructures for biodiversity data and observatories. Furthermore, it is referred to by the German research project SoKNOS³ delivering Service-Oriented ArchiteCtures Supporting Networks of Public Security as well as by the European project DEWS⁴. DEWS aims at developing an open, standard based Distant Early Warning System for the Indian Ocean, especially tailored to tsunami hazards.

The SensorSA is being applied in the SANY pilot applications, which are presented in chapters 4 and 8. The SensorSA and its major components are being reused in the German research project EWS Transport⁵ with the aim of developing an Early Earthquake Warning System that reduces the risk of damage for transport lines. Furthermore, the SensorSA is considered in the CEHIS final report (CEHIS, 2008) as enabling concept for the connectivity between environment and health information systems.

The objective of the SensorSA is to motivate and specify the basic design decisions derived from user requirements and generic architectural principles. Its focus is on a platform-neutral specification, i.e. it provides the basic concepts and their interrelationships (conceptual models) and abstract specifications.

¹ <http://www.orchestra.eu.org> or see the ORCHESTRA book edited by Klopfer and Kannellopoulos (2008)
² <http://lifewatch.eu>
³ <http://www.soknos.de>
⁴ <http://www.dews-online.org>
⁵ <http://www.ews-transport.de>

By abstract it is meant that the specification is independent of the specifics of a particular service platform. Such an abstract specification comprises service specifications, information models and interaction patterns between the major architectural components, as illustrated below:



The specification of the SensorSA is structured around the concept of architectural viewpoints of the Reference Model for Open Distributed Processing (ISO 10746-1, 1998).

The RM-ODP explicitly foresees an engineering step that maps solution types, such as information models, services and interfaces specified in information and service viewpoints, respectively, to distributed system technologies. We describe this mapping step in terms of engineering policies. These policies constitute architectural blueprints that enable a system engineer to specify implementation architectures according to given user requirements, as outlined in the lower part of this graphic.

SANY has performed this engineering step for the use cases which have been introduced in chapter 4; the resulting applications are described in more detail in chapter 8. Some generic patterns for such implementation architectures are described later on in the present chapter.

The SensorSA is a multi-style SOA. This means that, in addition to the classical architectural style, which is oriented towards remote invocations, it also supports an event-driven and a resource-oriented architectural style.

As such, it foresees mechanisms to generate events and distributes them as notifications to interested consumers. This enables spontaneous distribution of information about changing configurations in underlying sensor networks, e.g. the dynamic addition or removal of sensor devices, which is a pre-requisite for the support of the 'plug-and-measure' type of operation.

Furthermore, the SensorSA embeds a resource-oriented architectural style. Resource-orientation in the SensorSA refers to unique identification of geospatial resources (e.g. time series of observation results, spatial data sets) and their representations as tables, maps or diagrams. This approach provides more flexibility in the design of an implementation architecture, for instance, it enables the mapping to and the co-existence with so-called RESTful web service environments (Richardson and Ruby, 2007). By this multi-style approach, it remains a design decision of the system engineer in the engineering step which architectural style best suits the individual purpose and requirements.

This chapter provides an introduction to the design principles, the sensor model, the major architectural elements, the standards used and the service and interfaces that have been specified. In addition, it also presents data fusion methodologies and generic architectural patterns ranging from sensor networks, data fusion environments up to decision support infrastructures.

The complete specification of the SensorSA (Usländer (ed.), 2009) is available as public document and can be downloaded at the SANY project website.

DESIGN PRINCIPLES

6.1

A SOA for an open sensor-based environment cannot solely rely on existing design principles that are typically applied in commercial SOA environments (Erl, 2008). The SANY architecture team has refined them in the following way:

- **Rigorous Definition and Use of Concepts and Standards**
The SensorSA makes rigorous use of proven concepts and standards in order to decrease dependence on vendor-specific solutions. This helps

to ensure the openness of a sensor service network and support the evolutionary development process.

■ **Loosely Coupled Components**

The SensorSA allows the components involved in a sensor service network to be loosely coupled, in which case loose coupling implies the use of mediation to permit existing components to be interconnected without changes.

■ **Technology Independence**

The SensorSA is independent of technologies, their cycles and their changes, as far as practically feasible. Accordingly it is possible to accommodate changes in technology (e.g. lifecycle of middleware technology) without changing the SensorSA itself. The SensorSA is independent of specific implementation technologies (e.g. middleware, programming language, operating system).

■ **Evolutionary Development – Design for Change**

The SensorSA is designed to evolve, i.e. it shall be possible to develop and deploy the system in an evolutionary way. The SensorSA is able to cope with changes of user requirements, system requirements, organisational structures, information flows and information types in the source systems.

■ **Component Architecture Independence**

The SensorSA is designed in a way that service network and source systems (i.e. existing information systems, sensors and sensor networks) are architecturally decoupled. The SensorSA does not impose any architectural patterns on source systems for the purpose of having them collaborate in a service network, and no source system shall impose architectural patterns on a SensorSA. Important here is that a source system is seen as a black box, i.e. no assumptions about its inner structure are made when designing a service network.

■ **Generic Infrastructure**

The SensorSA services are independent of the application domain, i.e. they can be used across different thematic domains and in different organisational contexts. Ideally, any update of integrated components (e.g. sensors, applications, systems, ontologies) requires no or only little changes to the users of the SensorSA services.

REFERENCE MODEL FOR OPEN DISTRIBUTED PROCESSING 6.2

The conceptual foundation for the SensorSA has been the Reference Model for the ORCHESTRA Architecture (RM-OA).

The RM-OA provides a platform-neutral abstract specification of a geospatial service-oriented architecture that responds to the requirements of environmental risk management applications. It comprises generic architecture services and information models based on and extending existing OGC specifications.

Klopfer and Kannellopoulos, 2008

The design of the SensorSA follows the guidelines and viewpoints of the **ISO Reference Model for Open Distributed Processing** (ISO/IEC 10746-1:1998). However, since a SANY system has the characteristic of a loosely-coupled network of systems and services instead of a ‘distributed processing system based on interacting objects’ as presumed by RM-ODP, the RM-ODP concepts are not followed literally. The RM-ODP viewpoints are applied on a big scale to the structuring of ideas and the documentation of the SensorSA itself, and on a small scale to the description of the sensor model in order to capture the multi-fold facets of the term ‘sensor’:

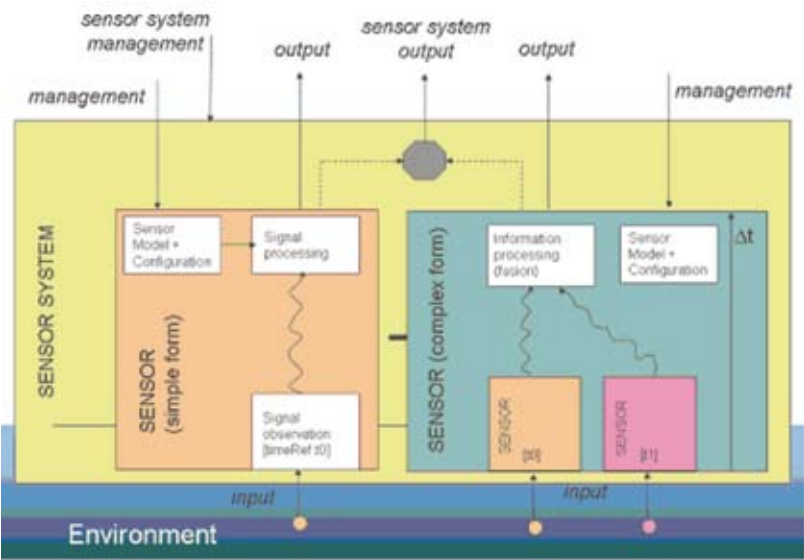
- The **Enterprise Viewpoint** of the SensorSA reflects the analysis phase in terms of the business contexts, related system and the user requirements expressed in use cases as well as the assessment of the current technological foundation for the SensorSA. It includes rules that govern actors and groups of actors, and their roles. Business examples are the European initiatives GMES, INSPIRE and SEIS and the world-wide initiative GEOSS. A use case example is the need to fuse earth observation products of GMES or GEOSS, e.g. optical images of a river estuary in a flooding situation with in-situ gauge observations of the river.
- The **Information Viewpoint** specifies the modelling approach of all categories of information, with which the SensorSA deals, including their thematic, spatial, and temporal characteristics, as well as their meta-information. Examples are information objects specified in class diagrams of the Unified Modeling Language (UML) and referred to by the specification of a fusion service.
- The **Service Viewpoint** specifies the interface and service types that aim at improving the syntactic and semantic interoperability between services,

- source systems and environmental applications. Examples are specifications of the externally visible behaviour of a service type, e.g. UML specification of the interface types of the fusion service.
- The **Technology Viewpoint** specifies the technological choices for the service platform, its characteristics and its operational issues, e.g. the specification of the platform ‘Web Services’ including a profile of the Sensor Model Language, or the physical characteristics of sensors and sensor networks.
 - Finally, the **Engineering Viewpoint** specifies the mapping of the service specifications and information models to the chosen service platform, considers the characteristics and principles for service networks, e.g. synchronous or asynchronous interaction patterns, and defines engineering policies, e.g. about access control and resource discovery.

6.3 SENSOR MODELS

The SensorSA defines in detail what is meant by the term ‘sensor’. First of all, it is related to the term ‘observed property’ that identifies or describes the phenomenon, which is being observed, or, applying the concise definition of the OGC Observations and Measurements model (Cox, 2007), the ‘phenomenon for which the observation result provides an estimate of its value’.

The SensorSA defines a sensor to be an entity that provides information about an observed property as output. A sensor uses a combination of physical, chemical or biological means in order to estimate the underlying observed property. Note that, basically, these means could be applied by electronic devices or by humans. In the former case, at the end of the measuring chain electronic devices produce signals to be processed. In the latter case, humans enter the observation results in a data acquisition system as a basis for further processing.



Furthermore, also simulation models or geo-statistical calculations are encompassed by this definition. They are then considered as a kind of ‘virtual’ sensor that could indeed replace or even complement sensor devices. With a view towards sensor devices the following sections take a look at different forms of a sensor.

From a technical point of view, we consider a sensor to be a device that responds to a (physical) stimulus in a distinctive manner, e.g. by producing a signal. This means that a sensor device converts the stimulus into an analogue or digital representation. Furthermore, we distinguish between simple and complex forms of sensors and sensor systems.

In its simple form a sensor observes an environmental property which may be a biological, chemical or physical property in the environment of a sensor, at a specific point in time (t_0) at a specific location, i.e. within a temporal and spatial context.

The location of the sensor might be different from the location of the observed property. This is the case for all remote-observing sensors, e.g. cameras, radar, etc. For an in-situ observing sensor, locations of sensor and observed property are identical, i.e. the sensor observes a property in its direct vicinity.

The simple form of a sensor provides information on a single observed property as illustrated in the left-hand sensor example of the illustration on the following page.

The observed property is usually converted to a different internal representation, usually electrical or mechanical, by the sensor. Any internal representation of the observed property is called a signal.

Within the sensor any kind of signal processing may take place. Signal processing typically includes linearization, calculations based on calibration coefficients, conversions to different representations and any calculations to prepare the sensor data for output.

A signal may also be transferred over longer distances. This could also be performed by a person carrying a chemical probe, e.g. a water probe from a river, to a laboratory. The path from signal observation to the output of signal processing takes time and may also be distributed across several locations. However, the temporal context (t0) and the spatial context of the signal observation must be preserved!

As an example, consider the above mentioned water probe measurement: It is imperative to preserve the time and the location at which the probe has been taken. Depending on the application context, the time and location of the examination of the chemical probe in the laboratory might be an essential part of the probe data, or it may be considered as additional meta-information. Finally, the observed property is accessible at the output of the sensor in a machine processable representation. The output provides information about the time (t0) and spatial context during observation, though those parameters are usually provided in the form of meta-information and not as part of the observation result. Due to the delay, Δt , produced by the sensor during the observation, the information at the output of the sensor cannot be accessed before $t0 + \Delta t$. This Δt can take any range from nanoseconds to several weeks or months.

Different sensors may provide different representations of the same observed property.

They may differ in the units, the quality of the representation, the observation method or the internal signal processing that was used. The estimate of the value of the observed property may be a single value, a range of values, a choice between worst and best value, a sequence of values or a multi-dimensional array of values representing, for example, a picture. It may contain values for each

point in spatial/temporal context or it may be a statistical representation in space or time.

The description of the representation as well as all other observation related information has to be provided as sensor meta-information at the sensor output to be used by an application.

A sensor may internally store representations of an older temporal context (history) or spatial context. In addition to its output, a sensor may provide an interface to perform the management of the sensor itself. For instance, this interface may be used to tag the sensor with a name, to configure the internal signal processing, or to monitor the behaviour of a device.

If an observed property cannot be observed with available simple sensor technology, it is possible to build a complex form of a sensor by using several simple ones, as illustrated on the right hand side of our illustration above. The information about the observed properties of the individual components of the complex form may be processed by any method of information processing (e.g. in fusion blocks). The output of the complex form of a sensor represents an observed property as defined by the sensor operator. This means that the linkage of the output of the complex form of a sensor to the output to the simple forms of a sensor is transparent. Still, even the complex form has to provide some information about the temporal and spatial context of its output data.

Several sensors may be combined to form a sensor system, which allows the management of the system that is holding the sensors in addition to the management of each individual sensor separately. This is done through the management interface of the sensor system. The key characteristic of a sensor system is its singular output and management interfaces that reflect its organizational unit.

The organizational unit varies in type and nature. Having a sensor system doesn't necessarily mean that the individual parts of the system do not provide individual interfaces. In addition, each part of a sensor system might be composed of sub-systems or individual sensors with individual interfaces as well. The key characteristic of the system remains its single output- and management interface, independently of any kind of interface provided in addition. Examples for sensor systems are satellites (whereas the physical structure of the satellite is a platform, not a sensor) with a number of remote-observing devices, weather stations with sensors for wind speed, temperature, and humidity, ground water observation systems used for surveillance of the environment around a chemical

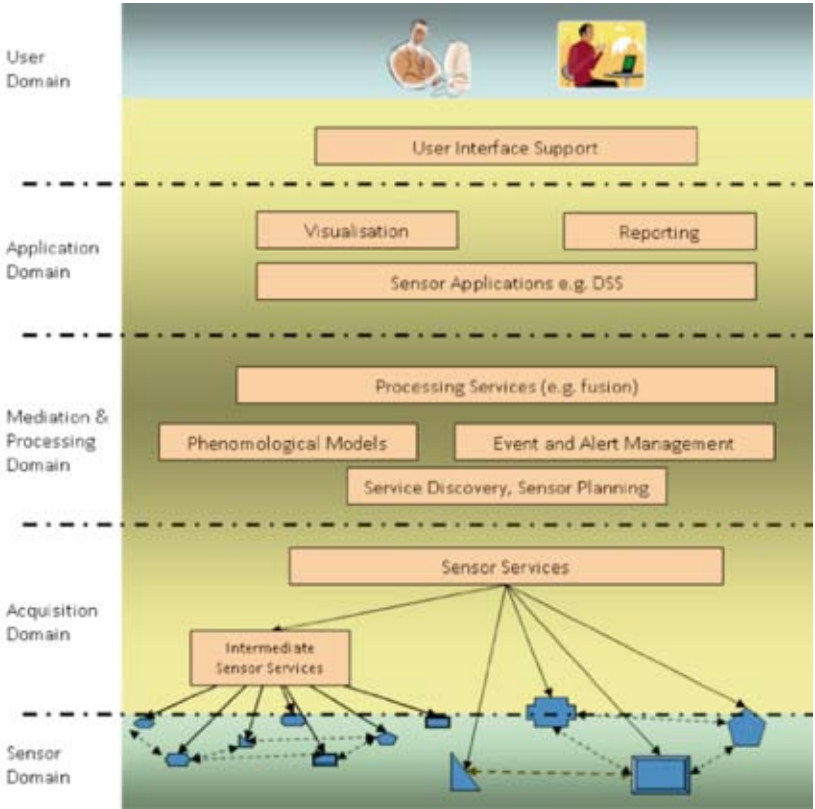
plant or a system of surface water observation points ordered on the surface and in the depth of a water body.

As opposed to a complex form of a sensor, the sensor system allows direct addressing of its individual parts as well as addressing of the sensor system as a unit. A complex form of a sensor provides only the management of the whole entity. Individual parts are not directly addressable. This difference affects the management interface, but has no influence on the response behaviour of both, complex form of a sensor as well as sensor system. Both might provide data that traces back to individual parts.

6.4 FUNCTIONAL DOMAINS

Services in the SensorSA are designed to support applications that serve the needs of users. They may call other services if this is required to fulfil the functions offered at their interfaces. In an extended situation, chains of service operation calls may be defined in order to realize more complex functionality. In a service network every service instance may call operations of any other service. The Service Viewpoint of the SensorSA categorises service types into functional domains expressing the area of concern for which they are basically designed:

- Services in the **Sensor Domain** cope with the configuration and management of individual sensors and their organization in sensor networks. They are abstractions from the proprietary mechanisms and protocols of sensor networks. An example is a take-over service in case of an imminent sensor battery failure.
- Services in the **Acquisition Domain** [AC] deal with access to observations gathered by sensors. This includes other components in a sensor network, e.g. a database or a model that may offer their information in the same way, i.e. as observations. The information acquisition process may be organised in a hierarchical fashion by means of intermediate instances, e.g. with data loggers.



- Services in the **Mediation and Processing Domain** [MP] are specified independently of the fact that the information may stem from a source system of type 'sensor'. They mediate access from the application domain to the underlying information sources. They provide generic or thematic processing capabilities such as fusion of information, the management of models and the access to model results. In addition, support for service discovery, sensor planning and the management of events and alerts are grouped in this domain.
- Services in the **Application Domain** [AD] support the rendering of information in the form of maps, diagrams and reports such that they may be presented to the user in the user domain.
- The functionality of the **User Domain** is to support the interface to the end user, typically in a graphical fashion.

6.5 SECURITY FRAMEWORK

Since the SensorSA aims to make web-based services easily accessible, access control and network security are an important issue to be considered across all functional domains.

Security in the SensorSA goes well beyond the usual scope of access control in distributed systems. It includes topics such as confidentiality and integrity of information, reliability of the sensor, service and network domain, protection of sensors, data sources and communication channels, as well as the traceability of workflows and the usage of resources.

However, a general solution for all of these problems and total security for sensor networks is beyond the scope of the SANY project and subsequently has not been addressed in full detail in the SensorSA security framework. Physical protection of hardware (deployed sensors), intrusion detection in source systems and protection against eavesdropping of communication channels require specific hardware, and application and situation dependent solutions.

As an open architecture, SensorSA does not specify what any particular sensor or service does to protect itself. What the SensorSA does include, are security provisions to control access to services that are considered part of the SensorSA. The focus of the Security Framework is on access control. In a nutshell, access to a particular service is controlled in accordance with a policy specified for that service.

This chapter outlines the major concepts of the SensorSA security framework. Let's dig a little deeper and see how this is accomplished.

6.5.1 Identities and Profiles

The first concept to be understood is Identity. There are three important constructs in the SensorSA identity model:

- **Identity**
An Identity is the basic entity in the authentication process. An individual subject who wishes to access a service must be authenticated as corresponding to a particular Identity. Collections of Identities can be organised as a Group.

- **Role**
Roles are an abstract concept capturing a set of Identities in terms of their function (e.g. 'administrator'). These are modelled as a special Identity attribute and can cross security domain scopes.
- **Group**
These are modelled as a special type of Identity, and are themselves composed of a set of Identities. In contrast to roles the scope of groups is limited to a single security domain.

These elements are encoded in tickets applying the Security Assertion Markup Language (SAML).

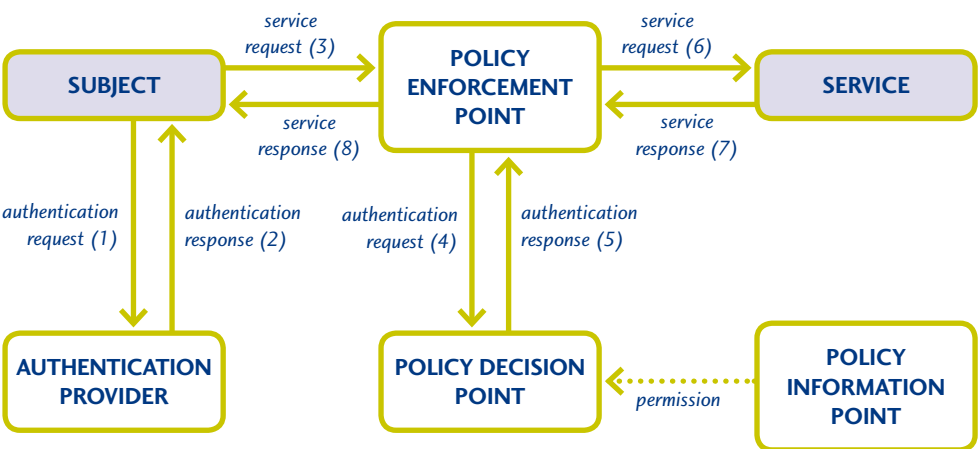
- Identity itself is a SAML Subject, and Identity Attributes are denoted in SAML AttributeStatements.
- Group Identities have no special type of SAML Subject, and are instead identified by the Attribute 'type' = 'user' or 'group'.
- Roles are identified by the Attribute 'role'.

There is a related concept called Profile. An acting entity maps onto a Profile, which itself can be related to a number of Identities. In this way an acting entity can use different Identities, each of which can be used for different purposes and verified with different methods. A Profile is composed of several profile attributes and is bound to one or more Identities. The profile attributes correspond to the properties of a user profile (name, organisation, email etc.) and follow a certain schema (e.g. the Lightweight Directory Access Protocol (LDAP) often used in Internet/Intranet application).

Access Control and Policy Enforcement

The second important concept in the security framework is that of a policy for access control. This policy specifies who may access a service and how it may be used. The fundamental steps in access control are as follows: Firstly, to authenticate the would-be users, i.e. to determine that they are who they claim to be, and then, secondly to determine whether they are authorised, according to the access control policy of the service, to access the service in the way they are requesting. This basic access control pattern is illustrated below:

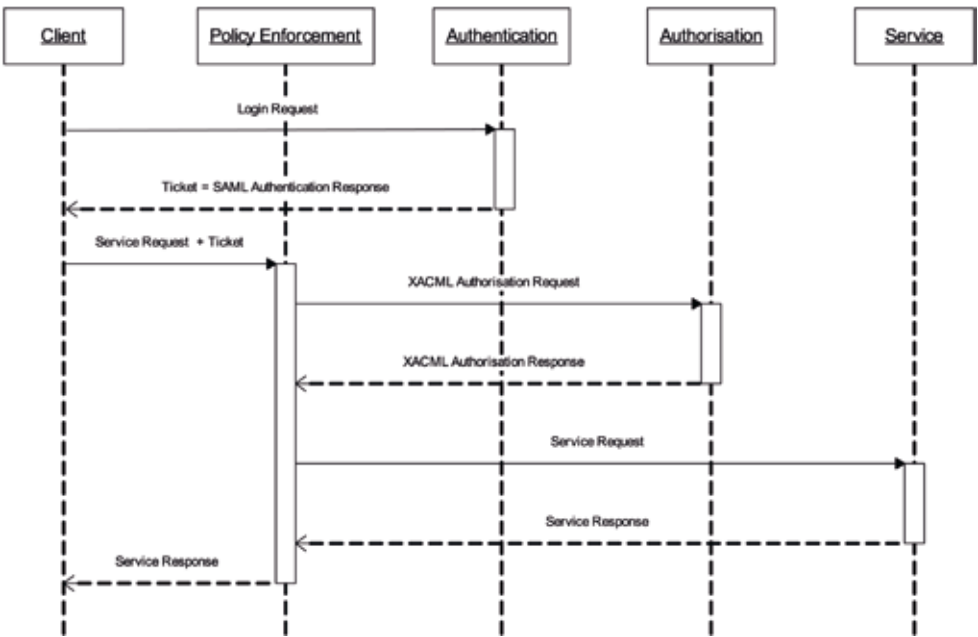
6.5.2



The elements in this pattern can best be understood by a typical transaction sequence:

- First, a Subject who wishes to use a service asserts its identity to an Authentication Provider. If the Authentication Provider determines the authenticity of the identity asserted by the Subject, it provides them with a validated 'ticket' (SAML Assertion, see below) with which the Subject may then issue a service request.
- Its service request goes to a Policy Enforcement Point. The Policy Enforcement Point verifies the ticket information with the Authentication Provider (a link not shown in the figure) and, after confirming an authenticated identity, requests authorisation of the service request for the Subject by the Policy Decision Point.
- The Policy Decision Point compares the request with the policy specification provided by the Policy Information Point. If it determines that the request is allowed, it issues a positive authorisation response to the Policy Enforcement Point. The rules for access control are specified in the XML dialect (Geo)XACML (see below).
- Finally, the Policy Enforcement Point delivers the service request to the service, passing the service response back to the Subject.

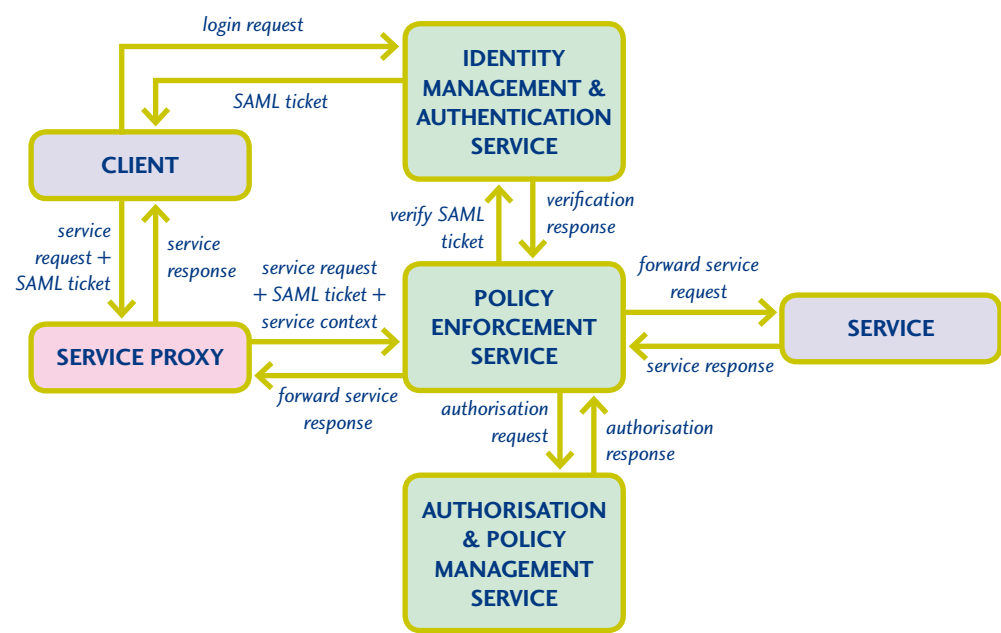
This sequence from top to bottom looks like this:



It's apparent that the components involved in the security framework effectively serve as a access control proxy layer for the Web service itself, fielding service requests in order to first authenticate the identity of the requestor, and then to authorise the requested access. This is key to preventing sensor access control from intruding on the services themselves. Indeed, by using a transparent proxy approach the services need not know that this level of security is being provided.

Security Framework Services 6.5.3

The components of the Security Framework are themselves realised as SensorSA services positioned within the mediation and processing domain. Their role and interactions in the access control pattern are illustrated below



- The **Identity Management & Authentication** Service is responsible for the management of identities, their authentication, and the management of credentials and issuing of sessions. An instance of the Identity Management and Authentication Service acts as both authentication provider and identity provider. The service supports the management of groups (of identities) as a special kind of identity.
- The **Policy Management and Authorisation** Service supports the management of policies, acting as policy administration point by allowing the management (select, create, update, delete) of (Geo)XACML policies, as well as policy information point. Moreover, as an instance of the authorisation service interface it acts as policy decision point by providing a decision on whether some identity (e.g. a user or a service) is authorised to access a certain resource.
- The **Policy Enforcement Service** handles the necessary interaction (authentication and authorisation) to obtain the required access control decision and is independent of the controlled service (generic).
- The **Service Proxy** mimics the controlled service and delegates the service request to the Policy Enforcement Service.
- In addition to the services supporting the Service Access Control Pattern the **Profile Management Service** manages profiles and their relations to identities.

Modes of Access Control

6.5.4

The concepts presented here constitute a standards-based mechanism for access control in service networks. They provide all that is necessary in order to equip service infrastructures with access control mechanisms with minimal effects on service interaction.

This includes a model for subject-related information, which can serve as a basis for access control across security domains, e.g. service networks of different stakeholders involved in a common environmental application. The information model supports several modes of access control:

- PBAC (Policy Based Access Control)
- IBAC (Identity Based Access Control)
- RBAC (Role Based Access Control) and
- ABAC (Attribute Based Access Control) which enable designers to cope with arbitrary requirements for the entity on which a decision is mounted.

The major advantage of the security framework is that service developers do not need to consider access control aspects when designing their services. Furthermore, the approach ensures backwards compatibility which means that an unsecured client can invoke service operations of a secured service and vice versa.

In compliance with the work performed in the OGC Security and Distributed Rights Management working groups, the SensorSA security framework incorporates prominent OASIS security standards, with the additional benefit of security aspects like message confidentiality and integrity that are already covered by the OASIS security standard family. Tangible results of this work are a set of tools (proxy generator, adapter template, administration interface) and a set of service implementations that can be used to secure arbitrary Web services.

6.6 THE STANDARDS STATE-OF-THE-ART

6.6.1 Standards Applicable to Conceptual Models

Strictly speaking, there is no such thing as a technology independent standard. Nevertheless, some of the standards relevant to environmental monitoring ICT infrastructure are sufficiently generic to remain valid, or in the worst case evolve in the course of the next 10-20 years, rather than becoming obsolete when the underlying technology changes.

In order to support the above mentioned design principles of ‘rigorous use of standards’, ‘technology independence’ or ‘generic infrastructure’, the SensorSA is based upon the following standards on the conceptual level:

- **The Reference Model for Open Distributed Processing** (ISO/IEC 10746 RM-ODP), which is used to structure the ideas and documentation in both the SANY and the ORCHESTRA Integrated Projects.
- **ISO 19101:2002 Geographic information -- Reference model**, which is a base for all OGC services.
- **Unified Modelling Language** (UML) as visual general purpose modelling language specified by the Object Management Group (OMG).
- **ISO/TS 19103:2005 ‘Geographic information -- Conceptual schema language’** provides rules and guidelines for the use of a conceptual schema language (here: UML) within the ISO geographic information standards.
- **ISO 6709:2008 ‘Standard representation of geographic point location by coordinates’** is applicable to the interchange of coordinates describing geographic point location. It specifies the representation of coordinates, including latitude and longitude, to be used in data interchange. It additionally specifies representation of horizontal point location using coordinate types other than latitude and longitude. It also specifies the representation of height and depth that can be associated with horizontal coordinates. Representation includes units of measure and coordinate order.
- **ISO 19107:2003 ‘Geographic information -- Spatial schema’** specifies conceptual schemas for describing the spatial characteristics of geographic features, and a set of spatial operations consistent with these schemas. It treats vector geometry and topology up to three dimensions. It defines standard spatial operations for use in access, query, management, processing, and data exchange of geographic information for spatial

(geometric and topological) objects of up to three topological dimensions embedded in coordinate spaces of up to three axes.

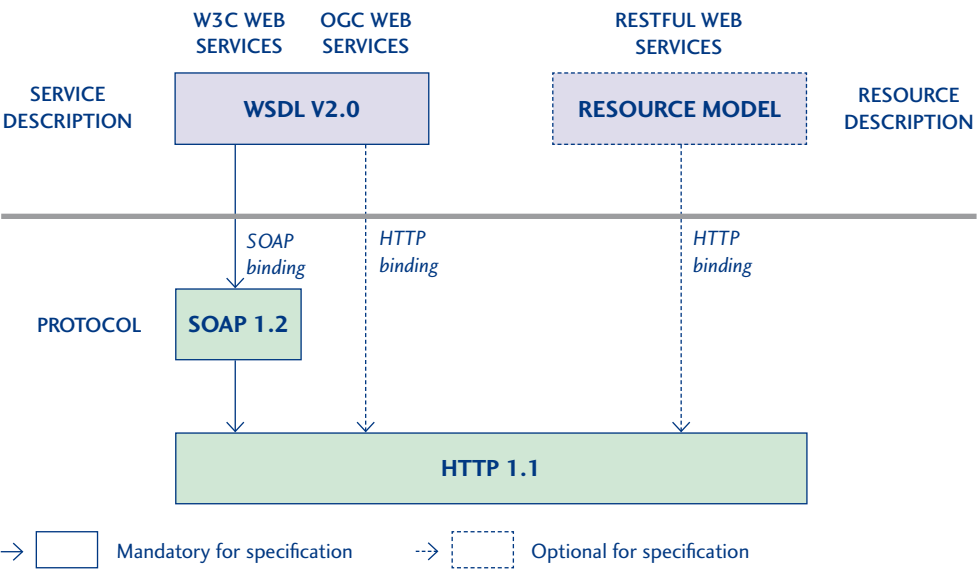
- **ISO 19108:2002 ‘Geographic information -- Temporal schema’** defines concepts for describing temporal characteristics of geographic information. It depends upon existing information technology standards for the interchange of temporal information. It provides a basis for defining temporal feature attributes, feature operations, and feature associations, and for defining the temporal aspects of metadata about geographic information. Since this International Standard is concerned with the temporal characteristics of geographic information as they are abstracted from the real world, it emphasizes valid time rather than transaction time.

Standards Applicable to Service Platforms

The SensorSA concepts are realised following the guidelines and technologies of standard (Web) service platforms. However, there are competing Web service paradigms on the market with disparate protocol bindings (e.g. SOAP or HTTP) and capability descriptions (e.g. service-oriented or resource-oriented).

6.6.2

In order to enable service interoperability, the SensorSA separates the platform specification into a core mandatory part and one or more optional parts as illustrated below. Three platforms are currently supported: W3C Web Services, OGC Web Services and so-called RESTful Web Services for the resource-oriented approach.



The core mandatory part refers to W3C Web Services which is, according to a decision of the OGC Technical Committee, also the strategic direction for all new specified OGC services. It requires a SOAP envelope embedded into an HTTP message as the transport protocol and WSDL as the service description language. Optionally, HTTP may also be used directly. This enables to also use the OGC Web services as they are specified today.

Furthermore, RESTful Web Services are supported. These rely upon the principle of Representational State Transfer (REST) (Fielding, 2000) which means, that the call of service operation is considered as a transfer of state information of uniquely identifiable resources in form of resource representations. In the SensorSA, the resources are typically geospatial resources described in a resource model (see below), e.g. a collection of sensor observations with a known geo-location reference, and their representations, which may be maps, tables or diagrams.

The multi-platform approach of the SensorSA facilitates the reuse and integration of existing software components and the evaluation of other service paradigms.

The technologies for these service platforms rely upon specifications of W3C⁶ (World Wide Web Consortium), OGC⁷ (Open Geospatial Consortium) and OASIS⁸ (Organization for the Advancement of Structured Information Standards). A short overview of these recommendations and standards and their importance for the SensorSA is given below.

6.6.2.1 Web Service Recommendations of OASIS and W3C

OASIS is a not-for-profit, international consortium that drives the development, convergence, and adoption of e-business standards. The OASIS Reference Model for Service Oriented Architecture (OASIS, 2006) specifies the common characteristics of SOAs independent of a particular service platform implementation. The SensorSA assumes these characteristics as requirements for service platforms to implement the SensorSA functionalities. The implementation technologies are provided by the W3C.

⁶ <http://www.w3.org>
⁷ <http://www.opengeospatial.org>
⁸ <http://www.oasis-open.org>

W3C develops interoperable technologies such as specifications, guidelines, software, and tools to realise the service platforms. The Web Services Architecture (W3C, 2004) identifies the functional components and defines the relationships among those components necessary to achieve the desired properties of the overall architecture.

W3C Web Services refer to distributed software systems designed to support interoperable machine-to-machine interaction over a network. They are built upon four main components (Fensel et al, 2007):

- an agreed transport protocol (usually HTTP),
- a platform-independent message description format (usually SOAP, see below),
- a language for Web service interface descriptions (WSDL, see below), and
- a registry for publication and discovery of available services (normally UDDI, but in an geospatial service platform the OGC Catalogue service is being used. It is described in this book in section 7.7).

For the message description format W3C proposes SOAP – a basic messaging framework specified as an XML schema that expresses the structure of request and response messages. Furthermore, it provides a standardised way how to handle faults. The message contents is conveyed in SOAP envelopes, typically using the W3C application layer protocol HTTP. As discussed above, the service platforms supported by the SensorSA combine these W3C standards in different ways.

Web Services Description Language (WSDL) is the XML language recommended by the W3C for the description of Web services (W3C, 2006). It provides specification of essential Web service components such as operations, their grouping into interfaces, the structure of related input and output messages as well as their mapping (binding) to an underlying transport protocol.

There is also ongoing research work in the field of semantic extensions of the Web (Semantic Web), which has already led to a series of basic W3C recommendations such as OWL. OWL is the W3C Web Ontology Language to define and instantiate ontologies with an increasing expressiveness according to the sub-variant of the language used (OWL Lite, OWL DL, OWL Full). The semantic extensions of the OGC Catalogue which are presented in section 7.7 rely upon ontologies typically defined in OWL DL.

Geospatial Standards of OGC and ISO

The SensorSA falls into the category of a geospatial service-oriented architecture, i.e., it deals with resources that have a reference to a location on the Earth. Thus, the specifications of the OGC but also the ISO 191xx series of

6.6.2.2

geomatics standards are highly relevant for the SensorSA. ISO 19109 provides a framework for geospatial information whereas ISO 19119 is dedicated to geospatial services, respectively.

Standardised OGC services provide the call of geospatial services, e.g. for the access to geospatial data sets (see the OGC Web Feature Service WFS, or ISO/CD 19142), the execution of geostatistical calculations (see the OGC Web Processing Service) and the generation of interactive maps from multiple geospatial servers (see the OGC Web Map Service WMS, or ISO 19128:2005). Furthermore, there is the OGC Catalogue service that facilitates the publication, the search and the discovery of geospatial resources.

An overview and a summary of the OGC approach is given in the version of 2008 of the OGC Reference Model (Percivall (ed.), 2008). Based upon these ISO/OGC standards mostly refer to the needs of the mediation and processing [MP] as well as the application domain [AD] of the SensorSA. Dedicated to the acquisition domain are the information models and services of the OGC arranged in the Sensor Web Enablement (SWE) Architecture (Simonis (ed.), 2008). They tackle the access to, the tasking and the management of sensors over the Internet and basically fall into two categories:

- OGC SWE encodings: a set of XML encodings for communication with sensors and access to sensors and other sensor-like information.

The most prominent SWE encoding standards are the OGC Observations and Measurements model (O&M) and the Sensor Modelling Language (SensorML).

- OGC SWE services: OGC Sensor Web Enablement working group developed a suite of service interface specifications used for communication with sensors and access to sensors and other sensor-like information.

The most prominent examples are the OGC Sensor Observation Service (SOS), the Sensor Planning Service (SPS) and the Sensor Alert Service (SAS).

Security Standards of OASIS and OGC

6.6.2.3

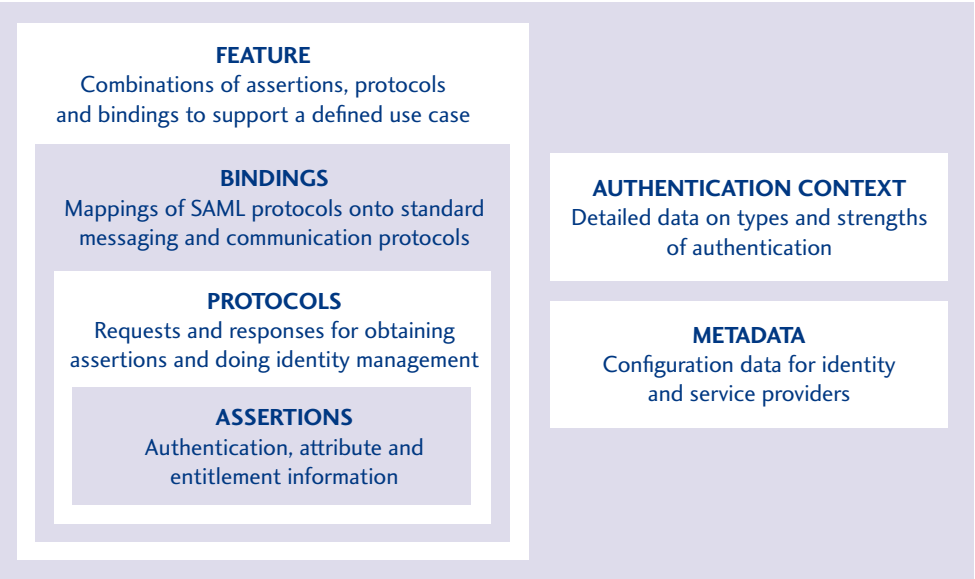
The SensorSA security framework uses two basic standards of OASIS:

- Security Assertion Markup Language (SAML)
- eXtensible Access Control Markup Language (XACML).

SAML is a language to encode security related information. In the SensorSA, especially SAML is used to encode identity-related information. SAML is summarised by (OASIS 2006) as follows:

SAML consists of building-block components (...) The components primarily permit transfer of identity, authentication, attribute, and authorization information between autonomous organizations that have an established trust relationship. The core SAML specification defines the structure and content of both assertions and protocol messages used to transfer this information.

SAML assertions carry statements about a principal that an asserting party claims to be true. The valid structure and contents of an assertion are defined by the SAML assertion XML schema. Assertions are usually created by an asserting party based on a request of some sort from a relying party, although under certain circumstances, the assertions can be delivered to a relying party in an unsolicited manner. SAML protocol messages are used to make the SAML-defined requests and return appropriate responses. The structure and contents of these messages are defined by the SAML-defined protocol XML schema.



Source: OASIS 2006

The means by which lower-level communication or messaging protocols (such as HTTP or SOAP) are used to transport SAML protocol messages between participants is defined by the SAML bindings. Next, SAML profiles are defined to satisfy a particular business use case, for example the Web Browser SSO profile. Profiles typically define constraints on the contents of SAML assertions, protocols, and bindings in order to solve the business use case in an interoperable fashion.

SAML core (assertion and protocol) is used exclusively in the SensorSA, i.e. no bindings or profiles have been defined.

XACML provides a policy language which allows administrators to define the access control requirements for their application resources.

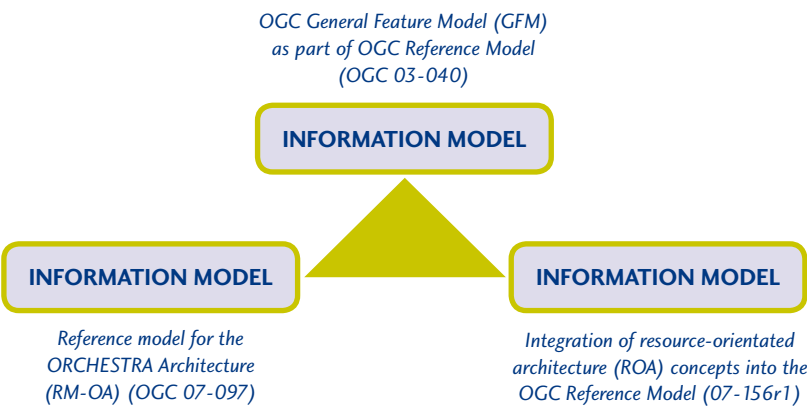
The language and schema support include data types, functions, and combinatorial logic which allow both simple and complex rules to be defined. XACML also includes an access decision language used to represent the runtime request for a resource. When a policy that protects a resource is located, functions compare attributes in the request against attributes contained in the policy rules, ultimately yielding a permit or deny decision.

GeoXACML is an extension to the OASIS XACML standard, which has been approved by the OGC. The primary goal of the GeoXACML extension is to support combinations of class-based, object-based and spatial permissions.

While class-based and object-based access control is already supported by XACML, the declaration and enforcement of spatial restrictions is not. GeoXACML defines spatial data types and spatial authorization decision functions, which can be used for additional spatial constraints for XACML based policies.

6.7 ELEMENTS OF THE SENSORSA

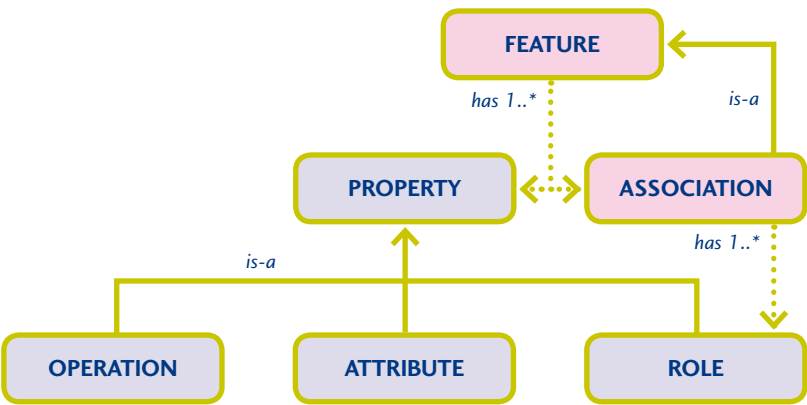
The SensorSA encompasses basic concepts such as the sensor model and the security framework described above, but also specifications of information and service models and engineering policies that provide guidelines how to use and combine these models. The models of the SensorSA are founded upon OGC standards, best-practices and submissions to OGC as shown below.



Let's now present an overview of these models and a selection of the related policies. Note that the models illustrated in this book are explanatory and do not have the rigor of the original UML models.

Information Models

The ultimate basis is the ISO/OGC-defined General Feature Model (GFM). The modelling unit of the GFM is the concept of a feature. Features play a very important role in the design of sensor-based applications as they represent entities in the universe of discourse of the users and stakeholders. In general, a feature is an abstraction of a real world phenomenon (e.g. a river or a forest). Features have properties which are usually attributes that describe thematic, spatial or temporal characteristics of a feature. Features may be associated to each other. This is expressed in terms of role properties of features as illustrated in the figure below.

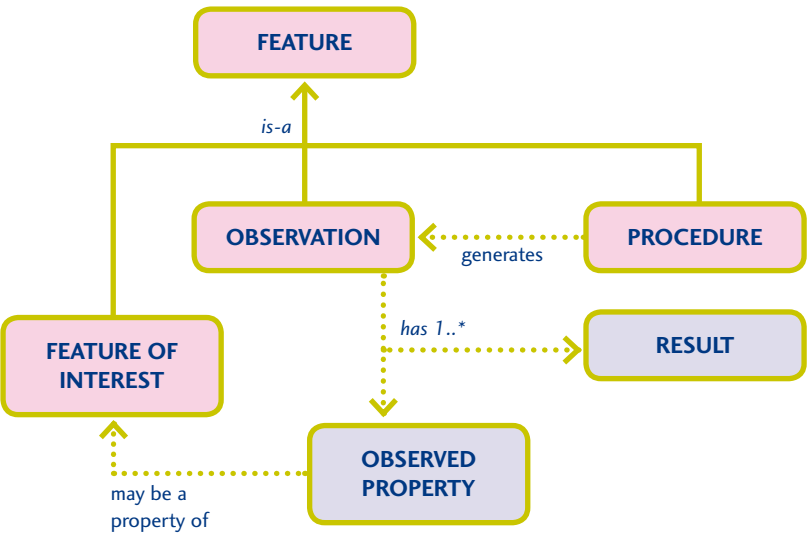


6.7.1

For instance, a feature ‘water body’ may be associated to another feature ‘gauge’ with the role ‘monitors’ on the gauge side and the role ‘is monitored by’ on the water body side. If required the act of ‘monitoring’ may itself be modelled as a feature in order to describe monitoring properties, e.g. to start/stop monitoring or to configure monitoring periods. A feature with a geospatial attribute, i.e. an attribute that describes a location relative to the Earth, is called a geographic feature. In sensor-based applications nearly all features are geographic features. They are the building blocks of project-specific application schemas, typically specified in UML and then mapped to XML in an engineering design step.

One extension of the GFM that is very relevant for the SensorSA is the OGC Observations and Measurement (O&M) model (Cox, 2007).

The O&M model is of core relevance for the access and interpretation of the data provided through the Sensor Observation Service.



The observation is the kernel concept. It is considered to be ‘an act associated with a discrete time instant or period through which a number, term or other symbol is assigned to a phenomenon’. The phenomenon is a property of an identifiable object, which is the feature of interest of the observation, i.e. the real-world object regarding which the observation is made.

The observation uses a procedure, which is often an instrument or sensor but may be a process chain, human observer, algorithm, computation or simulator. In the SensorSA the capabilities are defined in the Sensor Model Language (SensorML). The key idea is that the observation result is an estimate of the

value of some property of the feature of interest, and the other observation properties provide context or meta-information to support. An observation has the following characteristics:

An observation is modelled as a feature type whose instances are created at a specific time point or time period, the sampling time.

An observation may have been processed after sampling. The result time reflects the time when the result of the observation was produced.

The observed property identifies or describes the phenomenon for which the observation result provides an estimated value. It must be a property associated with the type of the feature of interest.

The procedure is the description of a process used to generate the result. It must be suitable for the observed property.

The result contains the value generated by the procedure. Note that the schema of the result data is not determined by the O&M model. The SensorSA recommends a self-describing schema, e.g. by using the definitions of the OGC SWECommon specification.

As further properties, an observation may have meta-information, e.g. the responsible actor for the observation and an indication of the event-specific quality.

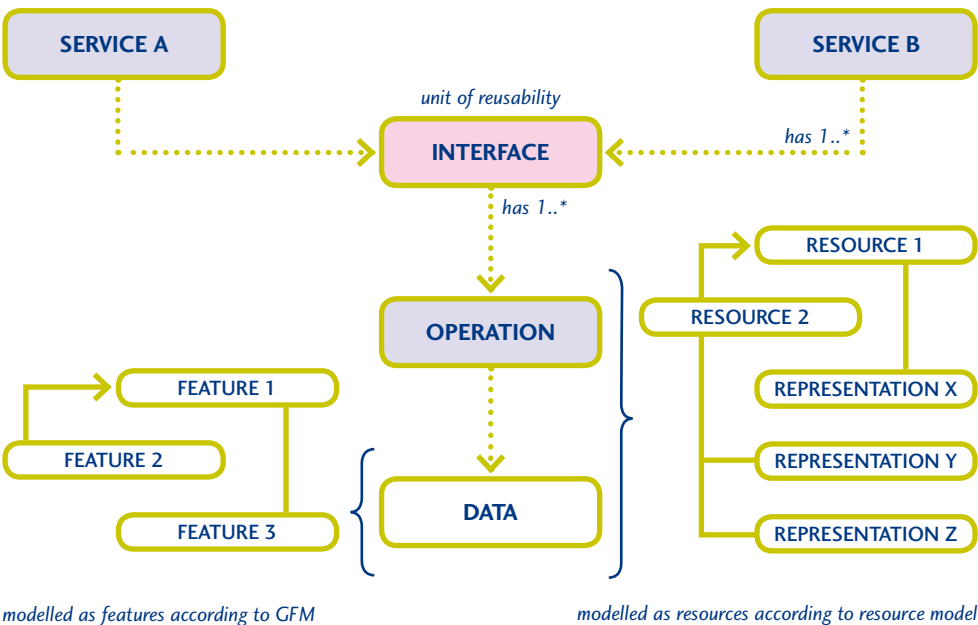
Service Models

6.7.2

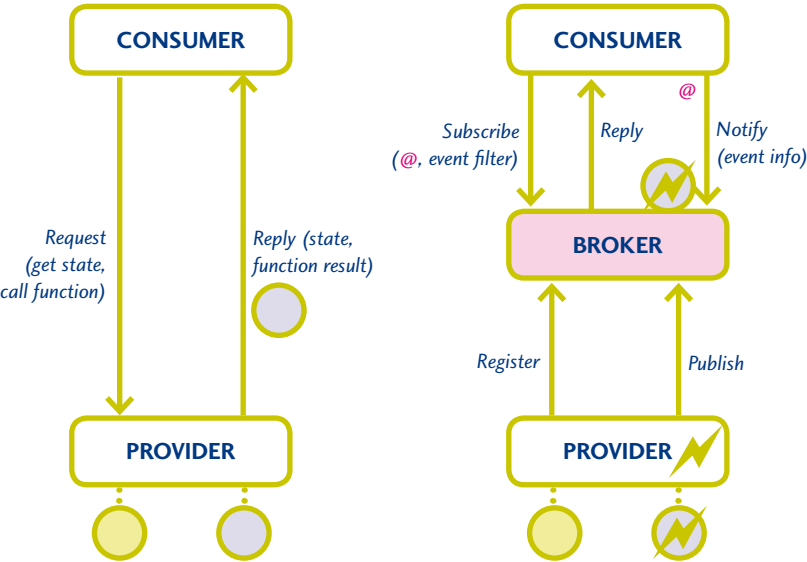
The SensorSA groups the services provided by a service platform into functional domains. For the specification of these services the SensorSA has adopted the service model of ORCHESTRA. This service model considers interfaces to be the unit of reusability on specification level whereby an interface is structured into operations. Operations access underlying data sets which are often related to attributes of features defined in application schemas according to the general feature model (see the left-hand side of the figure below).

Service types are specified in terms of one or more interfaces, whereby one interface may be attached to several service specifications.

For instance, the meta-information of services, their so-called capabilities, is specified in a dedicated capabilities interface which is common to all SensorSA services.



In the following table the currently supported service and interface types of the SensorSA are listed. In brackets you find the reference to the functional domain to which they belong. These services are oriented at the remote invocation and, partially, at the event-driven architectural style, e.g. the Sensor Alert Service and the OASIS Web Service Notification interfaces.



As shown above, the remote invocation follows a classical request/reply interaction pattern, e.g. to access the state of underlying data through a provider or call a function, whereby the consumer 'knows' the provider and waits until the reply has been received.

This is different in an event-driven interaction pattern. Here the consumer first declares interest in getting event notifications by issuing a subscribe operation with an event filter to a notification broker, usually together with a callback address (@). Then it continues with its activities. The provider publishes events to the broker, e.g. in case of a state change of an underlying resource.

The provider is unaware of the consumers that have subscribed to events. It is the task of the broker to analyse the event filters and to determine which consumer should be asynchronously informed by the broker about the event happening.

SERVICE/INTERFACE TYPE [FUNCTIONAL DOMAIN]	DESCRIPTION AND APPLICATION
Basic Interface Types [all]	■ Enable a common architectural approach for all architecture services, e.g. for the capabilities of service instances
Annotation Service [MP]	■ Relates textual terms to elements of an ontology (e.g. concepts, properties, instances).
Catalogue Service [MP]	■ Ability to publish, query and retrieve descriptive information (meta-information) for resources of any type. Extends the OGC Catalogue Service by additional interfaces for catalogue cascade management and ontology-based query expansion.
Feature Access Service [MP]	■ Selection, creation, update and deletion of features available in a service network. Corresponds to the OGC WFS but is extensible by schema mapping.
Identity Management- and Authentication Service [MP]	■ Creates and maintains identities. Supports the management of groups (of identities) as a special kind of identity. Proves the genuineness of identities using a set of given credentials and issues session information.

SERVICE/INTERFACE TYPE [FUNCTIONAL DOMAIN]	DESCRIPTION AND APPLICATION
Map and Diagram Service [AP]	<ul style="list-style-type: none">Enables geographic clients to interactively visualize geographic and statistical data in maps (such as the OGC Web Map Service) or diagrams.
Ontology Access Interface [MP]	<ul style="list-style-type: none">Supports the storage, retrieval, and deletion of ontologies as well as providing a high-level view on ontologies.
Policy Enforcement Service [MP]	<ul style="list-style-type: none">Handles authentication and sends authorisation requests to the Policy Decision Point for non-security enabled Web services.
Policy Management and Authorisation Service [MP]	<ul style="list-style-type: none">Provides a decision on whether some identity (e.g. a user or a service) is authorised to access a certain resource.
Profile Management Service [MP]	<ul style="list-style-type: none">Creates and maintains (user) profiles and their associations to identities.
User Management Service [MP]	<ul style="list-style-type: none">Creates and maintains subjects (users or software components) including groups (of principals) as a special kind of subjects.
Web Processing Service (WPS) [MP]	<ul style="list-style-type: none">Start, stop and result retrieval of information processes (e.g. statistical calculations).
Sensor Observation Service (SOS) [AC]	<ul style="list-style-type: none">Provides uniform access to observations from sensors and sensor systems that is consistent for all sensor types including remote, in-situ, fixed and mobile sensors.
Sensor Alert Service (SAS) [AC]	<ul style="list-style-type: none">Provides a means to register for and to receive sensor alert messages.
Sensor Planning Service (SPS) [AC]	<ul style="list-style-type: none">Provides a standard interface to task any kind of sensor to retrieve collection assets.
Web Notification Service (WNS) [MP]	<ul style="list-style-type: none">Service by which a client may conduct asynchronous dialogues (message interchanges) with one or more other services.

SERVICE/INTERFACE TYPE [FUNCTIONAL DOMAIN]	DESCRIPTION AND APPLICATION
OASIS Web Service Notification interfaces [MP]	<ul style="list-style-type: none">Family of related interfaces that define a standard Web services approach to notification using a topic-based publish/subscribe pattern.

Resource Model

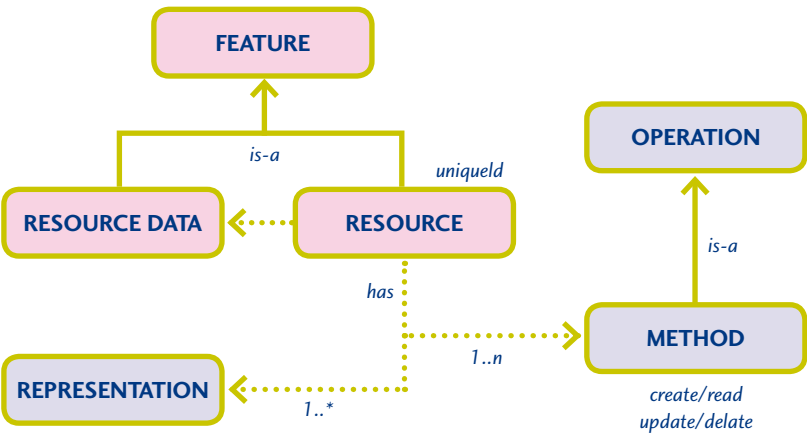
6.7.3

In addition, as mentioned before, the SensorSA also supports the resource-oriented architectural style. The right-hand side of the first figure in the last section illustrates the basic principle of this style. Instead of specifying dedicated operations for each functional need, we define resources with a unique identification (e.g. a URL in the Web) that provide a selected user-oriented view upon the underlying data.

As different users may have different needs, the resources may be retrieved in different representations, e.g. as a table, as a diagram or as a layer in a map.

Furthermore, in order to access and manipulate the state of the resources, there is a limited set of methods (operations) with a well-known meaning such as create, read, update and delete a resource.

The basic concepts of the SensorSA resource model as shown below are abstracted from the specification of RESTful Web Services according to (Ruby and Richardson, 2007).

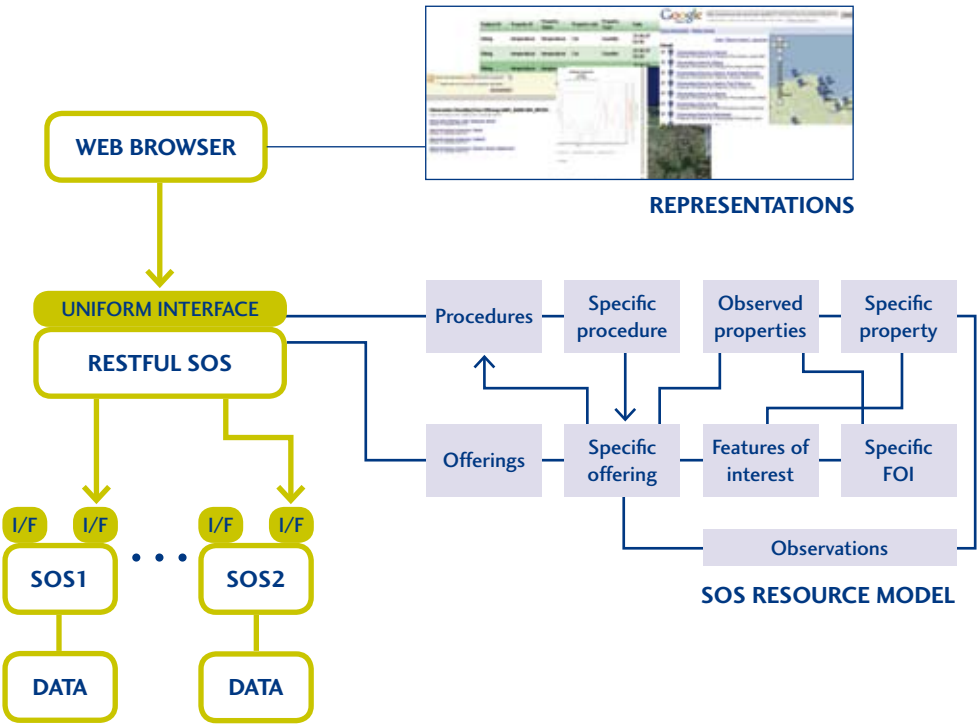


However, unlike the intense discussion in the Web service community, the SensorSA does recognise the resource-oriented architectural style as a beneficial complement to the remote invocation and event-driven architectural style.

RESTful Web services may be implemented on top of other SensorSA Web services to provide a simple user-oriented view for mash-up applications.

For instance, let's take a Sensor Observation Service (SOS). Here the major resources are instances of the O&M concepts such as features of interest, observed properties, procedures and observations and its capability concept of an offering. In combination they build a resource network through which a user may easily navigate using a Web browser and select the resource in which the user is interested in.

When retrieving the resource's state the user may then determine the representation form, e.g. by providing a well-known extension in the URL of the resource. This enables developers to easily embed sensor observation results into Web based applications, e.g. portals or Web sites of an environmental agency.



Processing of Quality Information

6.7.4

All data in SensorSA has an associated uncertainty depending on the available meta-information on how the data was observed (measured) or derived from other data sources. We first address measurement uncertainty and then uncertainty of general data.

Measurement uncertainties may be classified into two categories (ISO GUM 1993):

- Type A: uncertainty arising from a random effect; evaluated by statistical methods
- Type B: uncertainty arising from a systematic effect, evaluated by other methods

A common way of evaluating a type A uncertainty is to compute the standard deviation of the mean of a series of independent observations. A second common technique is an analysis of variance and random effects in data in dependence of experimental parameters.

Type B uncertainty is evaluated using scientific judgement. A typical cause is measurement bias due to the calibration of the measurement instrument or its behaviour in given environmental conditions (e.g. temperature, air pressure), or over time (deterioration of instrument, measurement drift). It is evaluated based on information about the instrument and environment. The measurement values may be corrected to compensate for known systematic effects.

Note the distinction between the terms error of a measurement and uncertainty.

Error is the difference between the measured value and the (in general unknown) 'true value' of the measured property.

Uncertainty is a quantified description of the doubt about the measurement result. The error of a measurement may be small, even though the uncertainty is large.

In SensorSA data arises not only from sensor measurements and observations, but also from data processing with specific services, e.g. a Kriging algorithm to generate a spatial coverage from a set of measurement points, or a time series analysis to produce a temporal interpolation. The results of such data processing steps are themselves uncertain, on the one hand due to the uncertainty of the input data, on the other hand due to the probabilistic or approximate nature of the processing itself.

Uncertainty of data is typically expressed with one of the following:

- Probability density function, e.g. a normal distribution with known mean and variance. The data value would then lie within one standard deviation of the mean with probability 68% and within two standard deviations with probability 95%.
- Intervals (the data value lies in [a,b]). This does not a-priori assume a uniform distribution on this interval; this would however be the case if the distribution of maximum entropy were chosen. An important special case is when then the measurement instrument can assert that the data value is below or above a given threshold, but can provide no further information.
- Statistics such as standard deviation and moments, or quantiles (the data value lies in [a,b] with probability 95%).

Within the SensorSA, the uncertainty of data sets is described using the UncertML (INTAMAP, 2007).

UncertML allows the information modeller to describe the uncertainty of a specific data set in an interchangeable way using an XML document conforming to the UncertML schema.

This XML document can be embedded in a SensorML document to express information about the uncertainty of some process. In addition, UncertML can also be embedded in an O&M document to express the uncertainty of a specific sensor observation.

6.7.5 Sensor Planning

Sensor planning in the SensorSA covers the aspects of sensor configuration (sensor tasking), sensor tasking feasibility analysis as well as updating and modifying sensor tasking instructions at runtime.

The goal of sensor planning is to hide the complexity of the sensor from the user. The same operation shall be provided to the user to task a buoy observing wave heights somewhere in the ocean, a simulation model calculating the weather for the next day, or a simple A plus B operation. The user shall only be confronted with a list of parameters that they might set (so called tasking parameters). All other complexity shall be hidden.

Sensor Planning takes place in each of the functional domains identified above:

- As an interface to the sensor domain, sensor planning allows (re-) configuration and managing of individual sensors, e.g. changing the sampling frequency.
- Sensor planning of the acquisition domain allows the tasking of individual missions. An example would be the tasking of a set of sensors that observe a specific area: a satellite with a mounted radar sensor, another satellite with electro-optical-sensors as well as some in-situ observations on ground are triggered to produce a complex data set of the area of interest.
- Sensor planning on the mediation and processing domain allows the integration of processing steps. Here, sensor planning may act as a process orchestration and chaining engine. A user might provide a set of interface locators that will be used to build a processing chain on the fly.
- The application domain as well as the user domain usually aggregate various sensor planning services and provide interfaces to the users. A user will be provided with a form that allows easy entry of tasking parameter data. These data are then sent to a sensor planning service on the application domain to execute necessary actions.

However, the same interface type is used to provide a façade to the tasking of each specific domain layer. This is achieved by the Sensor Planning Service (SPS).

Although the same operation (submit) is invoked for both planning and configuration the slight difference is the observation response. For planning the response encompasses observation data whereas the result returned upon configuration will contain the success status of the configuration step. One obvious advantage is the possibility of planning configuration tasks.

In general, sensor planning includes different interaction models or patterns. Some sensors allow synchronous interaction patterns, i.e. the service responds directly to incoming requests. An example would be an instance of an SPS that provides a facade for a simple forecasting model. This service, at least theoretically, could start unlimited parallel processes. Concurrent users don't compete for limited resources and the service can report the successful execution of the requested tasking right away.

Other sensors require asynchronous (event—driven) interaction patterns. This is the case if multiple users have to share a limited resource and the execution of the tasking cannot be handled instantaneously. An example would be a satellite that could at any moment in time observe a single scene only. If this satellite is equipped with an optical sensor, the observation depends,

among other factors, on the cloud coverage. Thus, the tasking request might consume any amount of time before being fully executed.

6.8 DATA FUSION AND MODELLING

Data fusion and modelling techniques are usually used to integrate observation data, contextual data and phenomenological models from different sources in order to obtain new environmental information where and when sensor measurements are not available. Observation sensors may include in-situ, air-borne and space-borne types, while models may include deterministic and stochastic models. In addition, data fusion numerical techniques provide a framework for integrating information uncertainties which are generated from sensor measurements and models with various inaccuracies. Several data fusion algorithms have been classified and developed in the SANY project. The classification exercise of these algorithms has been useful for hosting them under the SensorSA services and generically deploying them for multiple environmental risks and decision support pilot applications in SANY (Sabeur 2007).

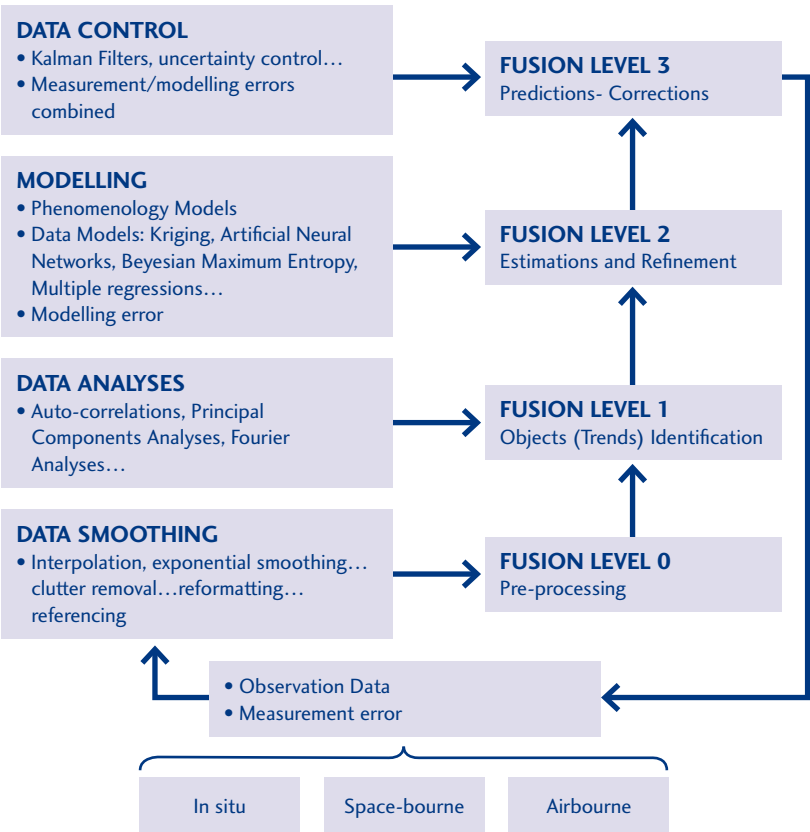
6.8.1 Fusion Levels

The generalisation of data fusion methods is the way forward for developing generic fusion services in the future. It will inevitably involve the classification of algorithms which specialise in the merger, correlation and modelling of data of different formats, spatial and temporal resolution and accuracies from various observation sensors. The sensors can be mobile or stationary.

The uncertainties on predicted parameters are the result of induced uncertainties from sensor measurement and those generated from numerical models.

Fusion techniques enable the predictions of environmental parameters and their respective uncertainties in time and space when or where sensing measurements are not available. Furthermore, those estimated uncertainties can be relatively decreased when new sensor measurements are obtained in time or sensors deployed in new areas.

A classification of fusion levels is provided below with illustrations of some of the typical numerical algorithms which are needed for processing observation data, identifying trends in data, modelling and controlling data with evaluated uncertainties:



The SANY requirements on data fusion and modelling led to the development of three distinct types of reusable fusion services:

- Spatial fusion services using Kriging or Bayesian Maximum Entropy
- Causal fusion services using multi-linear regressions or neural networks
- Temporal fusion services using state-space modelling and Kalman filters

These services have been implemented for environmental decision-support applications by various project partners in context with the SensorSA. They were then validated under multiple risk domain applications. These included pilot applications specialising in the prediction of microbial risk of exceedance in bathing waters in the Gulf of Gdansk (Poland), atmospheric pollution risks and false alarms in the City of Linz (Austria) and underground risks of subsidence in the City of Toulon (France), as outlined in section 4 and described in more detail in section 8.

Each of these types of fusion services are now presented in more details.

6.8.2 Spatial Fusion Services

Spatial data fusion services provide spatial trends of environmental parameters using observation data which are collated from a network of in situ sensors. This leads to the prediction of environmental parameters in areas where sensing is not available. The computation and analysis of spatial data uncertainties can also lead to identifying the areas where new sensor observations are required.

6.8.2.1 Kriging

Kriging is a method of spatial interpolation, which predicts values of an environmental parameter following observations of the same parameter at a finite number of sensors locations. The spatial predictions are simply weighted averages of the observed parameter values, according to the respective distances between the sensor points respective locations. The weights in Kriging are computed so that the variance is minimised. In this sense, Kriging is often called Optimal Interpolation.

The dependency of the interpolation weights on the distances between sensors is manifested in a variogram. The Kriging variogram essentially describes the variance of the difference between two distinct spatial observations. Furthermore, a realistic modelling of the variogram, should be based on reasonably accurate observations and a good understanding of the most dominant environmental processes that influence the spatial and temporal trends of the environmental parameter under study. This is of paramount importance for good Kriging results.

The numerical procedures in Kriging additionally involve the determination of measures of uncertainty when estimating environmental parameters in a spatial domain of interest. The approach leads to a good assessment of how observation sensors should be spatially distributed for achieving minimum uncertainty in spatial fusion.

Elevation correction

Since the meteorological stations which provide our wind data are located at different locations with different elevations above the sea-level we have implemented a pre-processing elevation correction step. We use wind profiling to transform wind vectors from the observation's elevation to the desired reference elevation. This is not required for the ground displacement pilot study.

Periodic variable support

The wind direction data was acquired in periodic directional formats as used in meteorology. This has required wind vector rotation to Cartesian references

prior to the stochastic analyses of data. About 80% of the wind direction angles span less than 180 degrees. In such cases, the creation of the variogram for Kriging requires the rotation of the wind vector in a way that the period onset does not intersect the wind directions span. This eradicates the periodicity issue in the data and enables the building of the variogram and proceed with data Kriging.

Ordinary Kriging algorithm

For theoretical-variogram model selection, eight models have been implemented in SANY. These include spherical, exponential, Gaussian, linear, power, generalised Bessel, sine hole-effect and cosine hole-effect. The model shapes are governed by a subset of the following parameters: nugget, range, power, hole and sill. Least-squares fitting methods are used to select a model with best fit of the experimental variogram. Background information about the phenomenon can be introduced by constraining the fitted model types and the parameter values and then, in effect, a variogram model reflecting the characteristics of the phenomenon of interest is finally selected. After selecting the theoretical-variogram model, model parameters optimisation is performed in order to improve the internal consistency of the model. Two statistics, termed Q1 and Q2, that need to be as close to their expected values as possible in order for the model to be consistent with the ordinary Kriging inductive bias have been used. Quadratic- sequential programming to tune the model parameters, subject to the parameter constraints discussed above, with a loss function proportional to the squared differences between Q1 and Q2 and their respective expectations we adopted. After the variogram model optimisation stage standard ordinary Kriging is performed and with mean and standard deviation computed.

Automated variogram selection

The ordinary Kriging procedure is combined with Automated Variogram Model Selection (AVMS). Background information describing the phenomenon characteristics can be reflected by the variogram model used for Kriging. For the ordinary Kriging with AVMS, along the sensor data, metadata is supplied that impose constraints to the variogram model to be selected in a way that reflects the phenomenology of the interpolated phenomenon. The most critical part of creating an experimental variogram is the selection of lags. Lags need to be selected so that they contain an optimal number of points in a way that the phenomenon physical characteristics are not smoothed out but that noise is not modelled. Generally the initial slope of the variogram needs to be well estimated so that the first few lags shall contain smaller number of points. If no hole-effect is expected the following lags may contain a large number of points,

but if the hole-effect is expected, then the lags shall contain a lower number of points and the effect is not smoothed out. The relative number of points in a lag is specified in the metadata supplied to the interpolation procedure. This relative number can be set by a phenomenon expert or automatically pulled from an expert system listing known phenomena.

6.8.2.2 Bayesian Maximum Entropy

Data fusion methods based upon Bayesian Maximum Entropy (BME) are able to consider soft sensor data, e.g. the sensor value lies in an interval, and additional phenomenological knowledge in the form of models. The results are statistics encompassing the uncertainty of the spatial/temporal interpolation given the uncertainty of the available information.

The overall BME fusion method is structured in three stages (Christakos 2000, Christakos et al 2002):

1. **prior stage:** consideration of general physical and scientific knowledge **G** about the spatio-temporal properties of the phenomenon of interest. This knowledge may, for example, be expressed in the form of spatio-temporal differential equations derived from physical laws or as covariance models. It is what is known before experience with the specific situation is applied. The prior probability distribution of the so-called random field of the phenomenon is determined using the maximum entropy (ME) principle, i.e. it is the most uninformative (unbiased) probability distribution given only **G**.

2. **meta-prior stage:** consideration of case-specific hard and soft data of the phenomenon of interest. This information is denoted by **S** (for specific knowledge) and is based on observations and measurements. Hard data refers to values believed to be accurate. Soft data is accompanied by uncertainty information such as a probability distribution for the value range.

3. **posterior stage:** processing (fusion) of the available knowledge **G** and **S** of the prior and meta-prior stages respectively to make a probabilistic map of the phenomenon for a given set of spatio-temporal points (typically a grid). The map is a statement of the general knowledge **G** relative to the case-specific knowledge **S** and is derived using Bayesian conditional probabilities.

If the general knowledge **G** comprises the mean and covariance, and if **S** includes only hard data, then the BME estimate coincides with the simple Kriging estimate (Christakos 2000, proposition 12.2). Similarly, if **G** is limited to the variogram and if **S** includes only hard data, then the BME estimate coincides with the ordinary Kriging estimate (Christakos 2000, proposition 12.3).

When applying the BME method in the SensorSA, the knowledge **S** is represented as an observation collection described with the O&M model and including uncertainty information in uncertML. The map resulting from the posterior stage is represented as coverage with associated uncertainty information.

Causal Fusion Services

6.8.3

Causal fusion refers to the indirect prediction of a target variable using a selection of explanatory variables. A number of causal fusion methods have been developed for use within the SANY project, including multiple linear regressions and neural networks. In most cases, historical target and explanatory variables along with real-time explanatory variables from in-situ sensors held in OGC compliant SOSs or from spatial fusion processes are accessed via an OGC compliant WPS. The resultant predictions are supplied to an OGC compliant SOS, and/or viewed through a web-interface.

Multiple Linear Regressions

6.8.3.1

Linear regression is used to construct a prediction formula for the target variable, given values of explanatory variables, by minimizing the sum of squared errors of linear fitting. Before constructing the linear regression formula, each explanatory variable is tested in order to determine whether a linear relationship to the target variable exists. The target variable is then predicted as a linear combination of the explanatory variables.

Linear regression is one of the most widely used modelling methods because of its effectiveness and completeness. Although the majority of processes are nonlinear in nature, many of them are well-approximated by linear models.

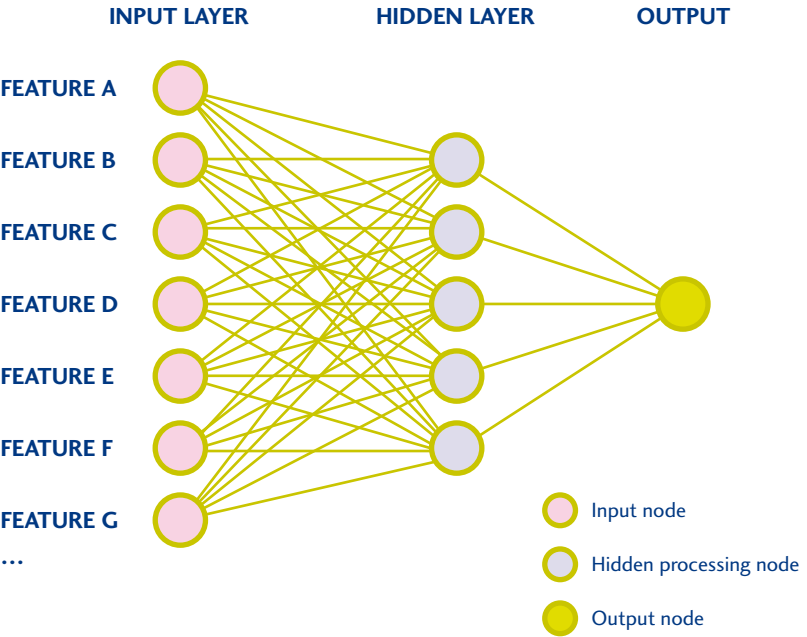
Linear regression enjoys solid theoretical background. The least squares criteria used for estimation of unknown parameters are optimal estimates under the most common assumptions for the model process. The algorithms are very efficient.

Linear regression estimates unknown parameters and assesses whether these parameters are statistically significant, which often has a clear meaning to scientific questions. Linear regression also assesses whether the model is statistically significant. The resulting model can be used to predict the target variable and confidence intervals.

6.8.3.2 Neural Networks

Neural networks are mathematical structures which are analogous to biological neural networks. The artificial neurons are set in layers and interconnected with each other. The neural networks are capable of processing non-linear statistical data and modelling complex relationships between inputs and outputs.

The most basic radial basis network consists of three separate layers. The input layer is the explanatory variables. The second layer is a hidden layer of high dimension. The output layer is the response of the network. The network topology is determined by the number of hidden units. One response is involved in this application. The neural network structure is illustrated in the following figure.



Neural networks are known for their ability to identify nonlinear relationships between explanatory and target variables. They have shown great prediction performance to fields where highly nonlinear processes are involved.

However, it is also generally considered a ‘black box’ approach since the model parameters are hard to interpret in terms of physical meanings.

6.8.4

Temporal Fusion Services

Temporal fusion can be used to predict the target variable directly from past observations of the target variable itself. The essential difference between temporal fusion and causal fusion is that temporal fusion takes the internal structure of data into account. In the SensorSA time-series data from in-situ sensors are obtained from SOS instances. The resultant predictions from the temporal fusion service are supplied to an OGC compliant SOS instance via a ‘virtual sensor’ controlled by an SPS instance.

Time-series analysis comprises methods to identify the nature of phenomenon in the sequence of observations and to make forecasts.

Methods for time series analysis are often divided into two domains: time-domain and frequency domain. The frequency domain approach is more suited to exploratory analysis. The time-domain approach is discussed here. Time series usually contains some typical patterns:

- Trend: represents long-run movements in the series.
- Seasonal cycle: seasonality pattern repeats itself more or less around a fixed period.
- Autoregressive component: represents data as a function of the past history plus a white noise.
- Moving average component: assumes data model is a linear combination of a prior random process.

Apart from the above regular patterns, an irregular component in the time series reflects non-systematic movements in the process.

The regular patterns can be identified through exploratory analysis or empirical knowledge of the process. At this stage, one must decide the order of trend, i.e., whether it is a random walk or a local linear trend, the existence of seasonal component and its period, the order of autoregressive and moving average components.

State-space Modelling

6.8.4.1

Once data patterns are identified, models for time series can be formed using an autoregressive integrated moving average model or state-space form.

The state-space form has enormous power to handle a wide range of time series models.

The basic structures such as trend and seasonal cycles are expressed explicitly in the model and are easy to interpret. The state-space form consists of a measurement equation and a transition equation. The transition equation contains the dynamics of the system under investigation and generates state variables. The measurement equation relates observable variables to state variables.

6.8.4.2 Kalman filters

After time series are modelled and put in state-space form, the Kalman filter algorithm may be used to produce predictions and smoothing of the state-space vector.

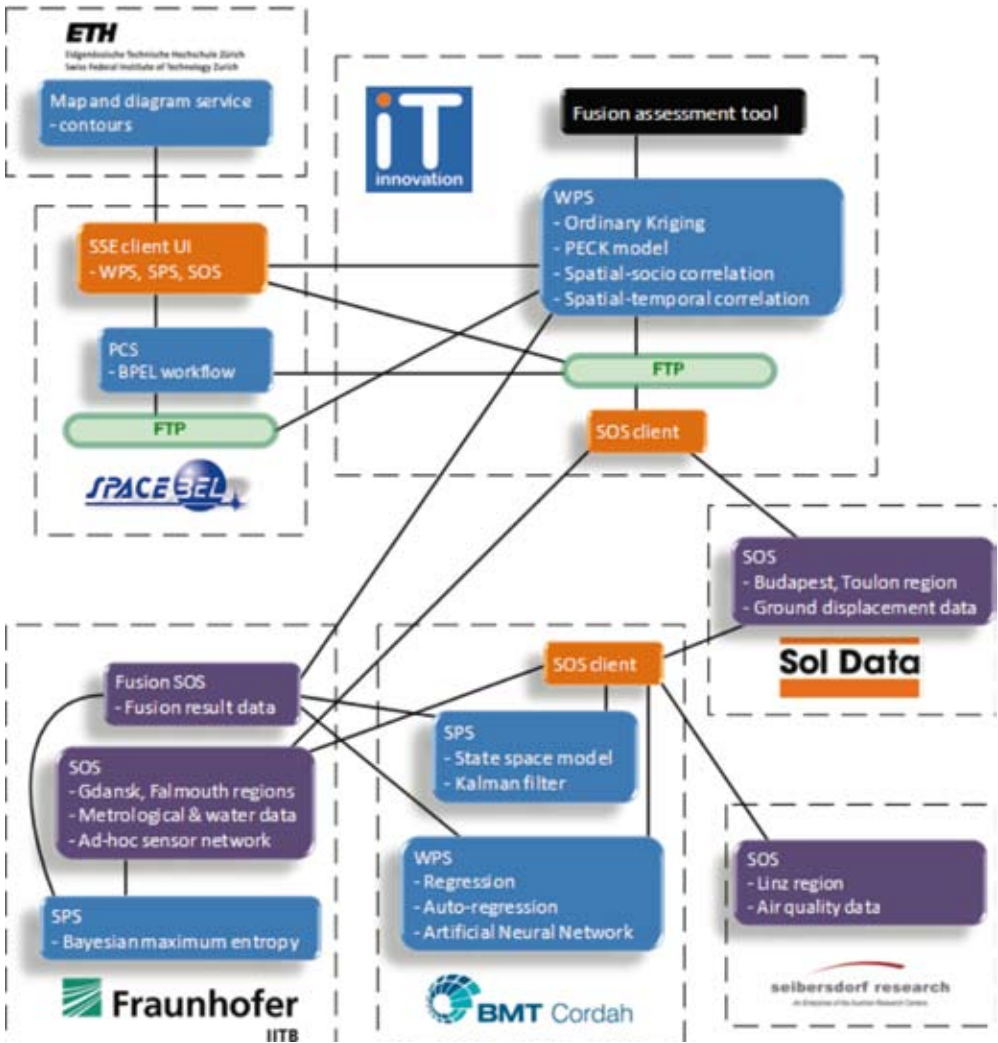
The Kalman filter is an important algorithm in many applications since it facilitates online estimation and enables the estimation and prediction of the state vector to be continually updated as new observations become available.

The Kalman filter is derived on the assumption that the disturbance and initial state vector are normally distributed. It gives optimal estimation of the state vector in the sense that it minimizes the mean square error within the class of linear estimators. It consists of two steps: prediction and update. The prediction step predicts the state variable and the prediction error to the next time step using the transition equation. The update step modifies the prediction once the observation at the current time step becomes available.

The Kalman filter also facilitates maximum likelihood estimation of the unknown parameters in the model. It enables the likelihood function to be calculated via prediction error decomposition. The maximum likelihood estimation can be carried out numerically or by an Expectation Maximization (EM) algorithm. The EM algorithm takes on a simple form comparing to the numerical solution and it always increases the likelihood during the iteration. The EM algorithm also tolerates missing observations and has a natural procedure to adjust the estimators.

6.9 IMPLEMENTATION ARCHITECTURES FOR FUSION SERVICES

There are various ways how fusion services may be implemented in an sensor service network based upon the SensorSA. Two examples of implementation architectures are presented in the following sections. The first one is using the built-in flexibility of the OGC compliant SOS/SPS, the second one embeds fusion into an OGC compliant Web Processing Service.



Fusion through Service Observation and Sensor Planning Services 6.9.1

A test bed has been developed for the fusion of sensor observations based upon BME (Kunz et al, 2009). The test bed implementation architecture has been designed for scalability and experiments in a wide range of scenarios, such as mobile sensors traversing several networks. At the sensor network level, an ad hoc ZigBee wireless network includes physical nodes that measure properties such as temperature, humidity, radiance and acceleration.

The objective of the test bed implementation architecture was to smoothly integrate a BME model into the landscape of OGC compliant sensor-related services: Incoming data is provided by SOS instances as sensor observations,

the configuration of the BME service is performed through an SPS instance, and, the output data is offered through a special SOS instance in the test bed, called a Fusion SOS.

One SOS server contains data originating from real sensors, whereas two other SOS servers handle data generated by a sensor simulator. The simulation facility allows experiments with many sensors of different types including tests with sensor data that is not uniformly distributed in space or time. Such data can be expected e.g. from sensors with intermittent availability or from moving sensors. The BME algorithms aim to fill the spatio-temporal gaps by computing intermediate values with an associated uncertainty depending on the quality of the input data.

This setup is complemented by simulated sensor nodes as illustrated in the figure on the next page.

The Fusion SOS Server generates its observations as spatio-temporal coverages using in this case a BME fusion algorithm that is parameterized and tasked by an SPS instance. An instance of the Map and Diagram Service is used to display the fusion results as a layer on a map.

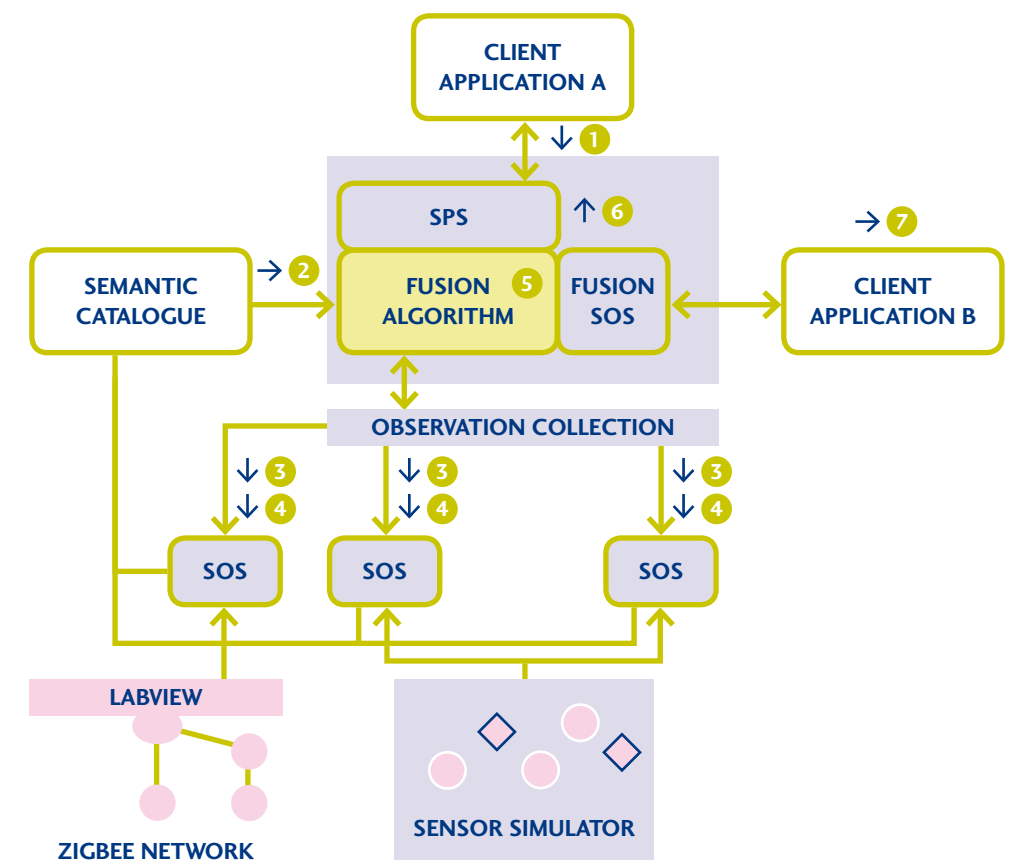
As a further component an instance of the catalogue service is being integrated, enriched by a semantically-enhanced query support. The Semantic Catalogue stores meta-information about all available sensors, services and observations and is used for the resource discovery in the test bed.

The overall fusion process flow comprises the following sequence of service operations:


1. A client application A wishes to create a new fusion result for observed property P in a time interval T and a set of sampling points S, e.g. a rectified grid, by accessing raw data from available SOS servers. This algorithm takes several configuration parameters as additional arguments and s described in SensorML for submission to the SPS. A prior *GetFeasibility* operation can be executed to check if the arguments are correct and acceptable. The SPS launches the fusion task. Its execution can take up to several minutes depending on the amount of data to be processed and the computational cost of the algorithm. The client may inquire about the execution progress with a *GetStatus* operation.
2. The fusion task queries the Semantic Catalogue for SOS servers with observations of property P in time interval T and in the area of a bounding box $BBox\{S\}$ around the sampling point set. In addition, the Catalogue

could have been queried in the previous step for suitable algorithms and SPS servers.

3. The fusion task applies the *GetObservations* operation to each SOS server to obtain the available observations of property P. Duplicates are recognized as observations taken by the same procedure (sensor) at the same sampling time; duplicates are deleted from the observation collection.
4. The fusion task determines the accuracy of the measurements. In the case of the test bed, this meta-information is in the SensorML of the related procedure. So the fusion task executes a *DescribeSensor* operation at the relevant SOS server to acquire this information. In general, the accuracy metadata could alternatively be in the observation result. The descriptive model language uncertML (INTAMAP, 2007) is used to encode the accuracy information into the XML file containing the result of the observation collection.



- 5. Now the fusion algorithm itself can be executed with the arguments
 - a) fusion parameters, b) the observation collection including (if available) the uncertainty of the observations, expressed as accuracy intervals, c) the sampling points at which the fusion is to estimate a value of the property. The result of the fusion algorithm is a coverage, i.e. a set of estimated property values for the sampling points together with a quantified description of their uncertainty. The uncertainty is described as a statistic such as variance or a probability distribution. The descriptive model language uncertML is used once again to encode the uncertainty information into the XML fusion result file.
- 6. The completion of execution of the fusion task is recorded by the SPS which can issue a notification to the client or another notification broker. The SPS responds to the operation *DescribeResultAccess* with the XML file argument required by a client when executing a *GetObservations* request to the Fusion SOS server to retrieve the fusion result.



Source: Fraunhofer IITB

- 7. Application Client B can, for example, display the fusion results geo-referenced and visualized using the SensorSA Map and Diagram service, in this case as a heat map.

Fusion through Web Processing Services 6.9.2

In the second example, a Kriging algorithm is made available through an instance of the OGC Web Processing Service (WPS). A WPS instance can either obtain sensor input data from a set of SOS instances or directly from client uploads to a local FTP site. Result data are provided immediately via a local FTP site and loaded into a fusion result SOS for archiving and persistent access as illustrated on the next page.

The WPS infrastructure supports generic fusion at two levels:

- 1. First a set of python scripts provides metadata driven pre-processing and post-processing.
- 2. Second a fusion framework provides configurable plug and play support for new algorithms, conversions and transformations.

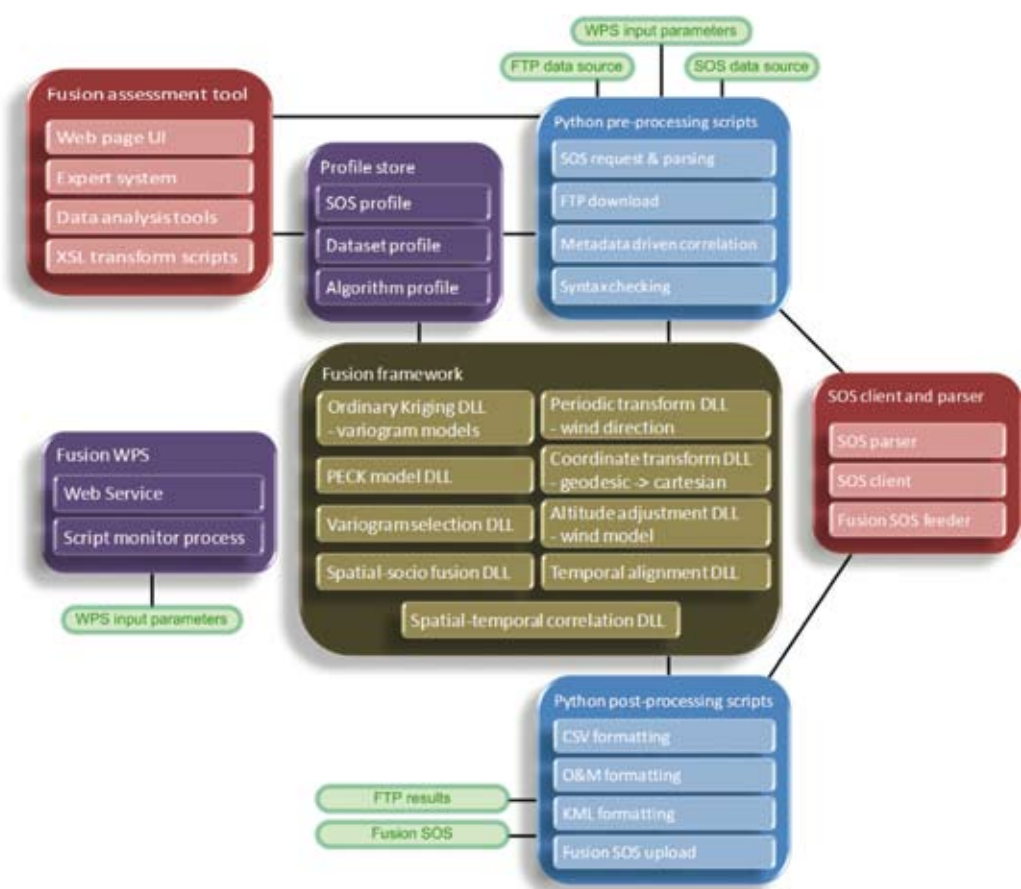
The pre- and post-processing python scripts provide metadata driven self-configuration based on the dynamically selected data source SOS/FTP and dataset provided by that data source. Each fusion observed target property, e.g. 'wind speed', is selected at runtime and the O&M metadata obtained from each SOS/FTP data source is used to identify the correct sensor value column(s), associated unit(s) and sensor accuracy information. Syntax checking is driven from the unit metadata provided. Post-processing scripts use the metadata provided in the input, e.g. using the same units when generating both CSV and O&M formatted result sets.

The python scripts themselves are designed to be either top level master scripts or generic sub-scripts. The top level master scripts are formulaic in design and can be automatically generated to semi-automate fusion deployment.

The fusion framework provides a simple framework with access to a bought in third party numerical library. The framework supports plug-in DLL's for various common mathematical tasks such as unit conversion, coordinate transformation and data fusion algorithms.

The idea is that over time new generic algorithms will be plugged in and the metadata and conversion/transformation services can be re-used.

The following illustration outlines the components of the spatial fusions service implementation:



- Discovery of sensor data and related services
- Access to sensor observations from different providers
- Management of sensor resources
- Subscription to and visualisation of sensor generated events and alarms
- Execution of processing services acting on sensor data
- Visualisation of sensor data on geographic map, charts, and tables

One of the key aspects in decision support are the fusion processing services with their ability to predict, in time and in space, the values of observed phenomena. Along with socio-economic data these predictions may also be used for impact assessment.

In the course of the SANY project, a web based multi-user Decision Support Infrastructure (Web portal) has been implemented to perform the above mentioned tasks.

The Web services used by the Decision Support Infrastructure are part of the SensorSA. The Decision Support Infrastructure provides a number of off-the-shelf clients for these Web services. Most of the clients are highly generic and instances of these clients can easily be deployed by a registered service provider on the portal.

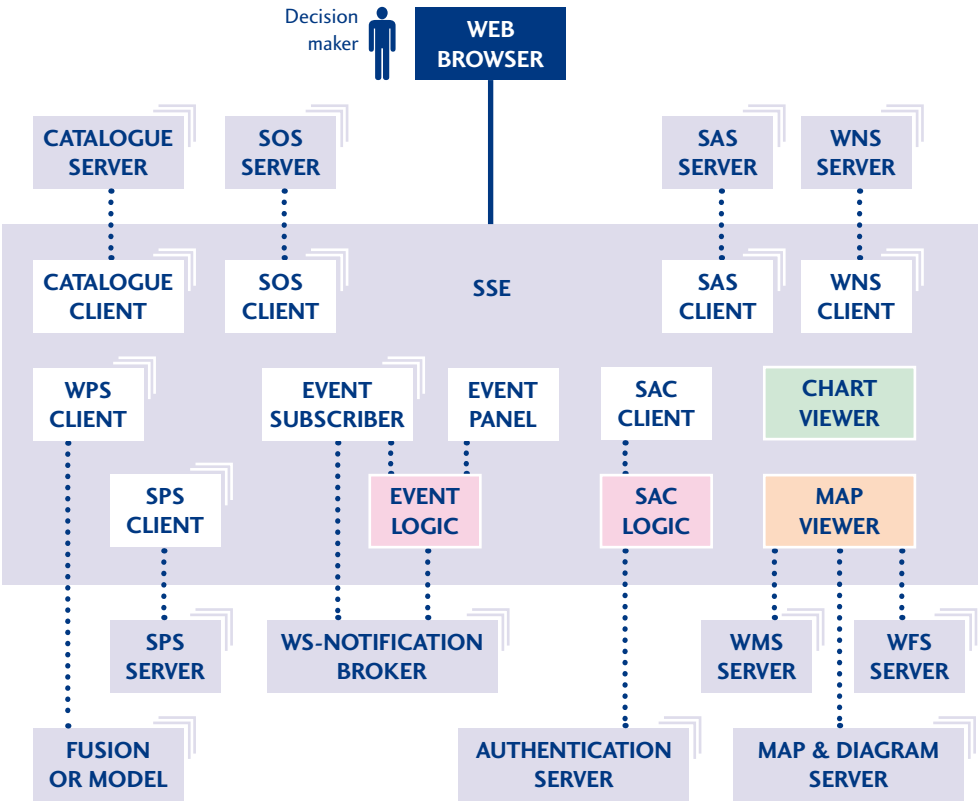
The generic usability of these clients is achieved by taking advantage of the service metadata that is available through the *GetCapabilities* operation and possibly other operations (e.g. *DescribeProcess*, *DescribeTasking*) exposed by these services in order to dynamically build the client input forms.

All the generic clients supporting OGC compliant SWE services (i.e. SOS, SPS, SAS, and WNS) can be configured to use SOAP instead of pure HTTP to communicate with the server (service instance).

The generic SOS client supports several result models: two standard specialized result models for time series and point spatial coverage and a more generic self-described observation model. The generic SOS client takes advantage of the SensorSA Map and Diagram Service to display contours on the map. For SOS service instances storing fusion results, the generic SOS client is able to display uncertainty information (expressed in UncertML) as well as sampling surface information for multi-point and rectified grid coverages.

6.10 IMPLEMENTATION ARCHITECTURE FOR DECISION SUPPORT

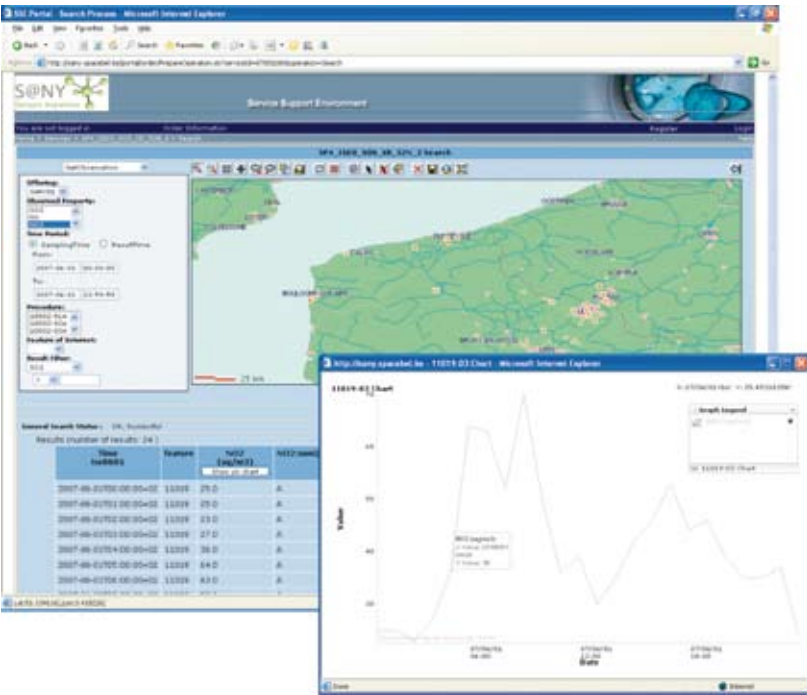
To help decision makers to assess and react to particular situations, an implementation architecture for a Decision Support Infrastructure has been designed based upon the SensorSA information models and services. It has the following main capabilities:



Clients are provided to subscribe to and receive events and alarms through various notification mechanisms: the OGC WNS service supports the notification to end users via a number of protocols (e.g. e-mail, SMS, etc) while the OASIS WS-Notification specifications support the notification to consumer services through an intermediate broker. The Decision Support Infrastructure includes a WS-Notification consumer that can be coupled with a WNS server and an Event Panel client to provide a very flexible notification infrastructure.

All the Decision Support Infrastructure clients can be configured to transparently support access to secured services i.e. services whose access is controlled by a Policy Enforcement Point according to the SANY security architecture. The clients automatically collect the assertion information for all the identities of the user (multi-domain security) through the SAC Logic which accesses the corresponding Authentication servers. This information (SAML tokens) is inserted in the SOAP header of all the service operations performed by the client. The user's identities are registered by him using the SAC client.

All the clients use the Map Viewer to graphically capture an area of interest on a map or display service results on the map. The illustration below shows a typical example of an SOS client that was deployed to monitor air quality in the area of Flanders (France and Belgium).



The implementation of the Decision Support Infrastructure is based on the ESA Service Support Environment (SSE) platform. The SSE portal provides many of the features needed to build distributed risk management applications based on open standards.

SSE was designed and implemented for the Ground Segment Department of the European Space Agency's Earth Observation Programme ESA EOPG, and continues to be extended to cope with new requirements and new interoperability standards. SSE initially allowed for the integration of Earth Observation (EO) and GIS services and data, but has now been extended (especially within SANY) to also include in-situ sensor services.

7 SANY Components

Building a SWE based system does not necessary imply that you have to do all the implementation work on your own. Instead you can rely on a broad spectrum of SWE component implementations, many of which are even available as Open Source Software.

In this chapter we introduce the major hardware and software components and concepts, which were successfully used within the SANY project.

7.1 SENSOR OBSERVATION SERVICE

The Sensor Observation Service (SOS) is the primary interface for accessing sensor data within the SANY architecture. The SOS is a standard Web service interface for requesting, filtering, and retrieving observations and sensor information. This is the intermediary between a client and an observation repository or near real-time sensor channel.

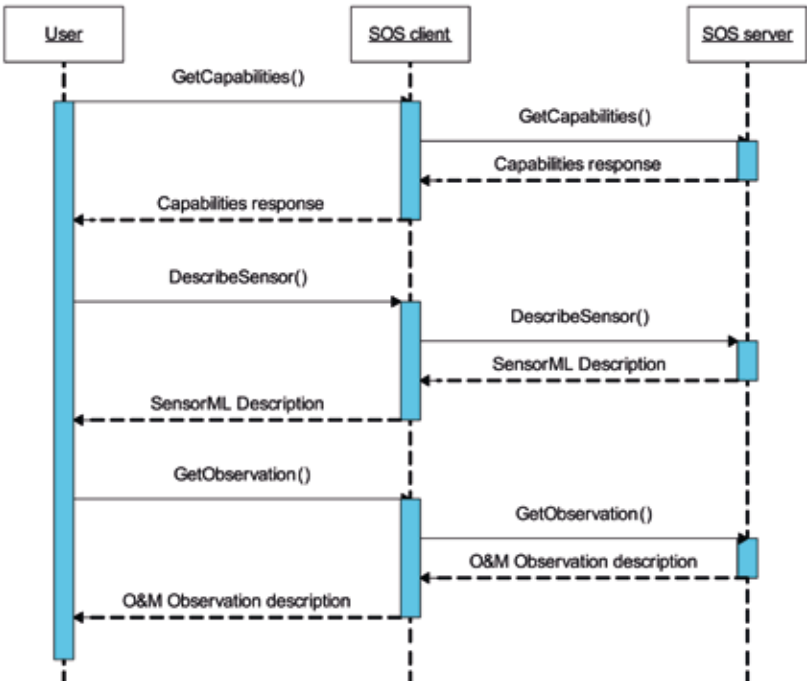
Within SANY, the SOS is used for Web enabling the sensor systems of in-situ sensor networks. The SOS is interrogated from the individual application services, e.g. a spatial fusion service, a settlement prediction model service, a temporal fusion service or from the SSE workflow, and is the key building block to facilitate the interrogation of sensors and visualisation of measurements.

The SOS implementations are based on the specification of the Open Geospatial Consortium, which comprises three different profiles:

- The **core profile** includes the following mandatory operations:
 - ◆ GetCapabilities for requesting a self-description of the service. It provides Service Provider information, the list of supported operations, and other information about the service.
 - ◆ GetObservation, for requesting O&M encoded sensor data, i.e. this operation actually sends back the observation data requested by the user.
 - ◆ DescribeSensor for requesting SensorML encoded metadata about the sensors contained in a SOS instance; the SensorML data which is sent back describes the arbitrarily detailed characteristics of the sensors and sensors systems used

- The **transactional profile** includes the following operations:
 - ◆ RegisterSensor for putting new sensors into the SOS
 - ◆ InsertObservation for inserting sensor observations
- The **enhanced profile** includes the following operations:
 - ◆ GetFeatureOfInterest, for requesting GML encoded representations of features of interest
 - ◆ GetResult for periodically polling sensor data
 - ◆ GetObservationByID for retrieving observations by passing the IDs of the observations

A typical SOS UML sequence diagram is presented below:



In the SANY implementation pilots, three Sensor Observation Service implementations have been deployed: a) an internal development of SolData, available under an Open Source license on the SANY website, b) the Open Source 52° North Sensor Observation Service and c) the Fusion SOS of Fraunhofer IITB, whose observations are coverages generated by a fusion process. All SOS implementations are of the 1.0.0 version of the OGC SOS standard. The formal OGC SOS compliance test is still to be established, but is currently in a beta phase and has been passed by the Fraunhofer Fusion SOS at this level.

7.2 CASCADING SOS

A Cascading SOS is a concept where an SOS service acts as the data source for an intermediate SOS, which itself provides a SOS interface to its clients. From an architectural point of view, using a cascading SOS may be of interest in a number of scenarios:

- optimisation of data flows
- provision of alternative data views
- (pre-)processing of data
- multi-level sensor data storage

7.2.1 Data Flow Optimization

On a conceptual level, data is directly accessible from the service provider or data source – on the engineering level, however, applications may face obstacles preventing efficient direct usage of an SOS by a client, such as:

- network performance problems
- limited resources on SOS servers
- different versions or feature sets of the SOS protocol in the client and server applications

Decoupling the data flow from the server to the client could include caching of data on the intermediate SOS service instance to overcome bottlenecks of limited or unstable networks or limited performance of the original SOS service.

7.2.2 Providing Alternative Views to Data

Using raw data as provided by a data source may not always be feasible or possible. Examples for such scenarios are:

- Different data providers may implement different data models for what is basically the same observed feature of interest.
- Data models used internally may not be feasible or appropriate for publishing them or making them available for a specific purpose.
- Organisations may need to provide an aggregated view of data collected by different providers, e.g. for implementing federated data pools.

The Cascading SOS concept can offer a solution to the requirement of offering an alternative view on it data sources by implementing an intermediate SOS server that provides a single interface to the underlying data sources. This

results in a clean distinction between the data access and processing on the client side, and the aggregation, transformation and/or filtering of the data that is necessary for a specific purpose in the intermediate SOS.

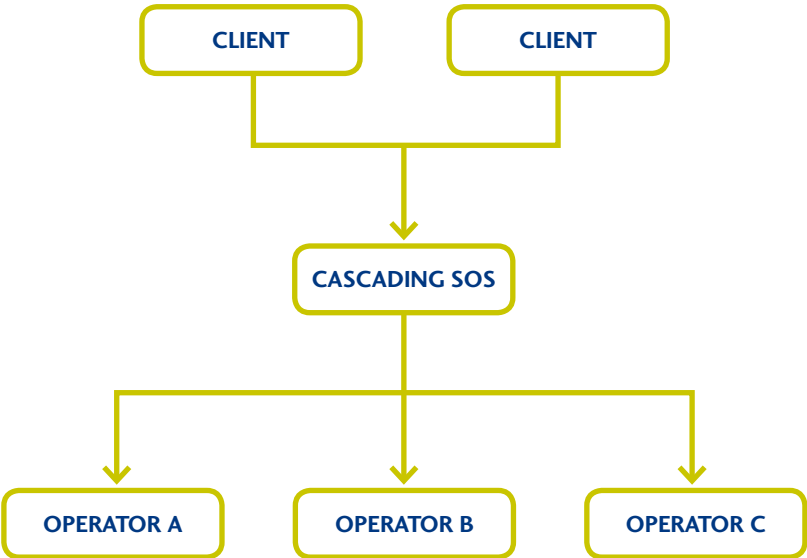
Data (Pre-)Processing

Depending on the client applications requirements there may be a need to pre-process data on the fly. This could be due to of limited computing capacity on the client side, e.g. when using smart phone based applications, or because data is requested in a predefined format that does not comply with the native data source. A typical simple scenario for data processing on the fly would be the calculation of mean values for time series data. While the measured data may be available with, for example, half-hour mean values from the sensors, an application may require daily mean values for its operation. In this case a cascading SOS could calculate the daily mean values on the fly and provide the cumulated results to the client application.

Multi-Level Sensor Data Storage

Sensors or data loggers connected to the sensors are often located in remote locations near the place where the observations are taken and tend to have strict constraints regarding storage space, which imposes problems for long time storage of observations.

To overcome these limitations, a cascading SOS can be used as illustrated below:



7.2.3

7.2.4

The cascading SOS fetches and stores the data provided by the sensors or data loggers, and the client applications access this service instead of accessing the sensors directly.

The prototype implementation of the Cascading SOS developed in SANY is based on TS-Toolbox (see section 7.10), and uses the ‘Formula 3’ for data (pre-) processing. This prototype can be also be used as a gateway to proprietary UWEDAT environmental data acquisition system, and to data stored in comma- or tab- separated data files.

7.3 SENSOR PLANNING SERVICE

The Sensor Planning Service (SPS) is the standard interface for all sensor, model and process tasking operations, whereby the latter two can also be handled with the Web Processing Service. SANY uses the Open Source 52° North SPS implementation as well as the Fraunhofer Fusion SPS that tasks fusion processes. Both are of the 1.0.0 version of the OGC SPS standard. The formal OGC SPS compliance test is still to be established, but is currently in the beta phase and has been passed by the Fraunhofer Fusion SPS at this level.

The following operations are specified within the SPS standard:

- **GetCapabilities** for requesting a self-description of the service
- **DescribeTasking** for requesting information that is needed for preparing a valid task, e.g. information about the necessary parameters.
- **GetFeasibility** for checking if a task with certain parameters can be executed or not, e.g. if the sensor is busy it might not be possible to successfully submit a task.
- **Submit** for sending a task that shall be executed by a sensor to the SPS.
- **GetStatus** for checking the status of a task, e.g. completed, cancelled.
- **Update** for updating the parameters of a task.
- **Cancel** for cancelling a task.
- **DescribeResultAccess** for retrieving information where the results of a task, e.g. the observations, can be accessed.

In addition to the operations specified by the OGC, the 52° North SPS offers additional functionality, which allows the administration of SPS instances. This includes:

- Registration of new sensor plug-ins and instances
- Unregistering sensor plug-ins and instances

- Updating a registered sensor
- Getting detailed status descriptions of sensor instances
- Updating information about the services providing access to the data collected by a sensor instance

The modular plug-in architecture allows the flexible integration of any kind of sensor data into an SPS instance. It offers an open, well-documented interface that can be used for easily developing plug-ins for connecting new sensors to your SPS. Through its plug-in concept, the integration of new sensors is fully supported and offers the flexibility to adapt an implementation to the specific requirements of your use case.

WEB NOTIFICATION SERVICE

7.4

The Web Notification Service (WNS) is mainly used to support asynchronous communication patterns, where message-originating services have to deliver messages to clients on protocols other than HTTP.

Since SANY mainly analyzed alternative asynchronous interaction patterns using WS-Addressing and WS-Notification standards, the relevance of WNS as a message-indirection service for SANY was rather low. However, since the functionality of a protocol transducer is required in many use cases, we have included its description in this chapter.

The WNS used in SANY is an Open Source implementation done by 52° North and based on the OGC WNS Best Practice Paper version 0.0.9. It includes the following set of operations defined in the WNS specification:

- **GetCapabilities** for requesting a self-description of the service
- **Register** for allowing clients to register themselves to the WNS by proving information about their communication endpoint (e.g. their email address). The registration of single users as well as of user groups is supported.
- **Unregister** for removing a client from the WNS
- **UpdateSingleUserRegistration** for allowing a client to provide a new communication endpoint (e.g. a new email address)
- **UpdateMultiUserRegistration** for adding or deleting members from a registered group
- **DoNotification** for submitting a message to the WNS, which will be forwarded to the specified receiver

The WNS allows the integration of a broad range of communication means for sending notifications. Adding a new communication channel just requires the implementation of a new handler, which forms the bridge between the WNS business logic and the communication system. By default XMPP and SMTP support are available. Furthermore a handler is available which supports sending SMS, fax and phone messages via the commercial HTTP-to phone/fax/SMS service 'ecall.ch'.

Thus, the available Web Notification Service implementation already provides a broad range of initially supported communication protocols and a flexible architecture for easily integrating additional notification channels, which even qualifies the WNS to be used within applications setups based on OASIS standard 'WS-Notification'.

7.5 WEB SERVICES NOTIFICATION

Due to the increasing demand for more flexible and dynamic services, communication patterns are required that effectively allow for the decoupling between the notification publisher and subscriber.

The Web Services Notification (WSN) aims to standardise the way in which Web services interact by using 'Notifications' or 'Events':

These specifications provide a standardized way for a Web service, or other entity, to disseminate information to a set of other Web services, without having to have prior knowledge of these other Web Services. They can be thought of as defining 'Publish/Subscribe for Web services'.

These specifications have many applications, for example in the arenas of system or device management, or in commercial applications such as electronic trading. (OASIS Web Services Notification (WSN) TC 9)

Another approach called WS-Eventing10 has been followed by the W3C, but it is not a W3C recommendation yet.

A high level overview of the WS-N functionality is provided by Niblett and Graham (2005).

9 www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsn#overview
10 Harmonization of the OASIS and W3C specifications was intended. Unfortunately, these efforts have ceased.

SENSOR ALERT SERVICE

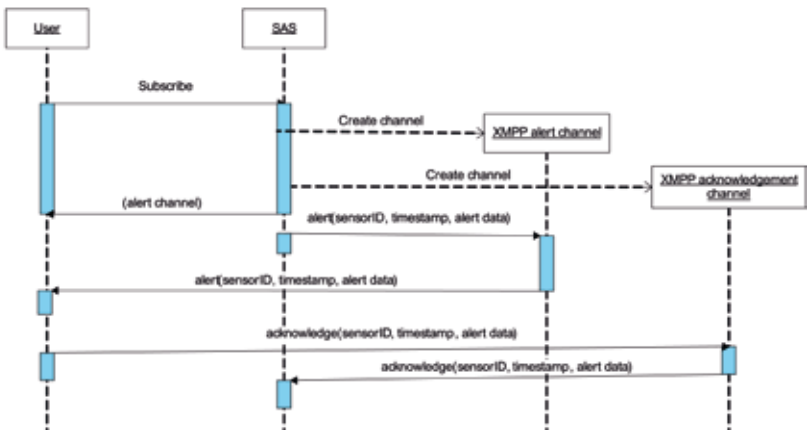
7.6

The Sensor Alert Service (SAS) defines an interface that allows nodes to advertise and publish observational data or alerts and corresponding metadata respectively. It also allows clients to subscribe to this data – or any other data that is produced by the SAS based on incoming messages from sensors – within specific thresholds. This observational data might be a single observation result, a complex observation result or even an *alert* in its nature. Within SANY the SAS has been used by end-users to subscribe to alerts and set alert conditions for the sensors of their choice, provided those sensors publish events through the SAS.

The SANY Sensor Alert Service implemenation is composed of four components:

- SAS server;
- SAS client;
- Multi User Chat (MUC) program;
- Jabber server that deals with XMPP messages.

Two protocols are used for the communication between the sensor, the server and the client. The sensors use SOAP over HTTP to advertise themselves, and the XMPP protocol to publish their data. The client sends the subscriptions and receives the answer on SOAP over HTTP, and receives alert messages from the client on XMPP. The SAS uses the Extensible Messaging and Presence protocol (XMPP) to provide the push-based notification functionality, used for instant messaging. Communication between the MUC and the client or server is done over XMPP. The Web service SOAP bindings are document literal with a wrapped parameter style. The SAS UML sequence diagram is shown below:



According to the OGC SAS Best Practice Paper version 0.9, the following operations are currently defined:

- **GetCapabilities** for requesting a self-description of the service

Sensors advertise to the SAS the data they publish. In return they receive the information where (to which multi user chat) they can publish their data. Thus, three operations for managing such advertisements are implemented:

- **Advertise** for allowing sensors to inform a SAS about the data they publish and returning the information where they can publish their data to
- **CancelAdvertisement** for cancelling an advertisement
- **RenewAdvertisement** for renewing an advertisement (in order to avoid that the advertisement expires)

Finally, three operations are available, which allow clients to subscribe to the information they are interested in and for managing these subscriptions

- **Subscribe** for allowing clients to subscribe to the information they want to receive
- **CancelSubscription** for cancelling a subscription
- **RenewSubscription** for renewing a subscription (in order to avoid that the subscription expires)

When subscribing to certain information at a SAS you are able to use the filtering options defined in the SAS specification. This comprises

- **Spatial filtering**, within a bounding box or at a certain feature.
- **Sensor based filtering**, i.e. by sensor id.
- **Content based filtering**, i.e. smaller than, greater than, equal to, not equal to.

The table below illustrates, which filtering enhancements are currently discussed for future revisions of a SAS specification:

	SAS	SASNEW
Spatial	(✓) only bounding boxes r points)	✓
Temporal	✗	✓
Comparing	(✓) not 'equal or less' nd 'equal or more)	✓
aggregation of conditions	✗	✓

It is expected that future developments like more powerful filtering capabilities will quickly be incorporated into SAS implementation.

CATALOGUE SERVICE

7.7

Catalogue services play an important role for the discovery of resources. Conventional catalogues usually contain meta-information about available data and service resources.

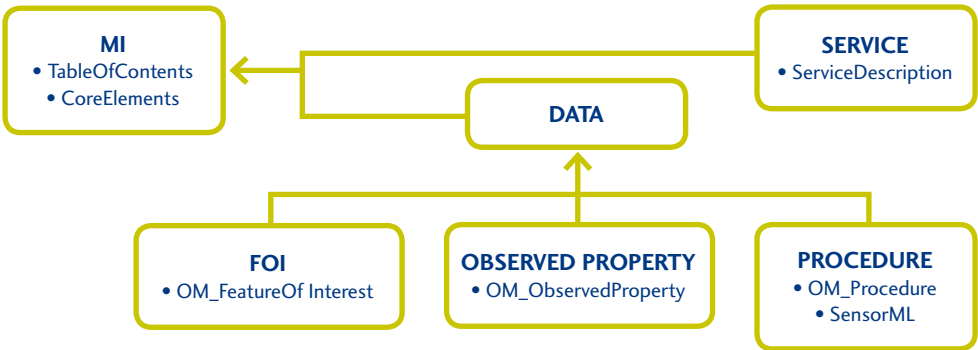
A typical user query to a conventional catalogue could include 'give me all services supporting standard interface x' or 'give me all datasets in a specific region, where the responsible party is y'.

A catalogue used for the discovery of sensor related meta-information needs to address additional requirements. Typical queries for such a catalogue differ from the conventional ones. Some examples may be: 'give me all 'temperature' observations in Austria of May 2009' or 'give me all entries supporting a specific sensor type'. Looking at these queries it is clear that additional search criteria and specific meta-information are needed, which reflect the needs from the sensor domain.

SANY addressed these challenges in developing a meta-information schema for the catalogue which follows the Observation and Measurement Model (O&M) from the OGC (Cox). This model is used by Sensor Observation Services, which provide the meta-information necessary to answer the queries above. Besides conventional catalogue resource types (data and service) SANY defined the following new meta-information resource types according to the O&M Model for the catalogue:

- The 'Feature of Interest' representing the observation target
- The 'Observed Property' describing the phenomenon to be observed (e.g. temperature)
- The 'Procedure' representing a specific sensor, sensor system(s) or algorithm(s) used by a system.

The illustration below shows the resource types of the so called SANY Application Schema for Meta-information. Each resource type supports mandatory meta-information sections (table of contents and core elements) containing common meta-information elements like 'title', 'keywords' or 'source url'. Further meta-information can be provided by specifying optional sections. This has been used for the description of the new resource types. Additionally the 'Procedure' resource type supports a SensorML section for a detailed description of sensors. The following figure shows the resource types with the mandatory and optional sections:



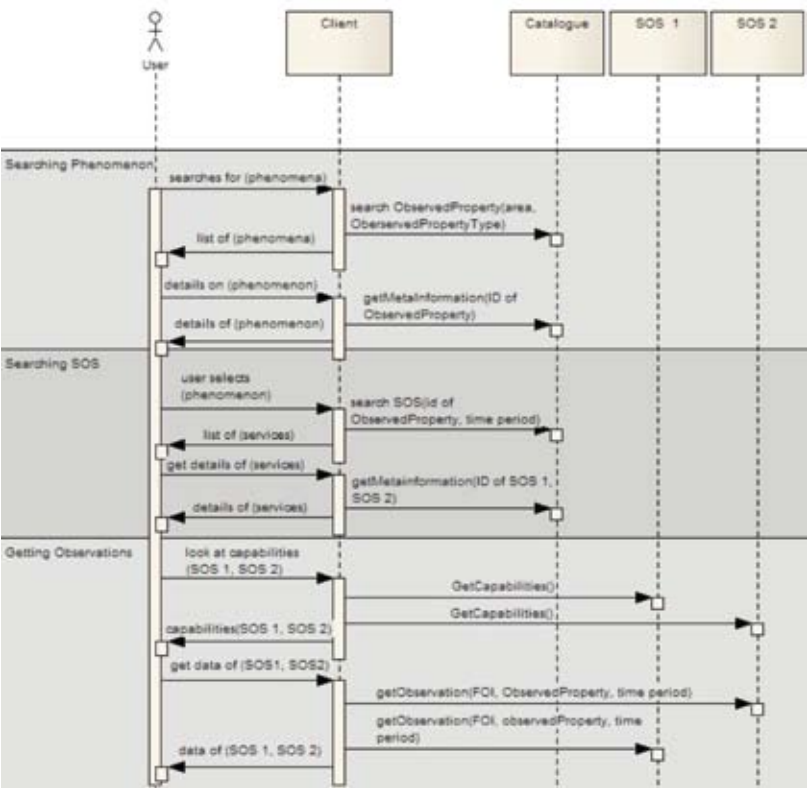
To support the possibilities of the SANY meta-information schema for the discovery process new search criteria so called 'queryables' for the catalogue were necessary. The following queryables have been defined:

- 'FeatureOfInterest' supporting the possibility to search for specific feature of interests, like a specific test region.
- 'ObservedProperty' supporting the possibility to search for general phenomena, like 'urn:ogc:phenomenon:temperature'.
- 'Procedure' supporting the possibility to search for general sensor types (e.g. accelerometer) and sensor instances.
- 'DatasetType' supporting the possibility to search for specific resource types like 'Feature of Interest', 'Observed Property' or 'Procedure'.

But how can the meta-information schema and the new queryables be used for the discovery of observations? The solution is a combined usage of the catalogue service and the SOS:

- In a first step a user can search for phenomena in the catalogue.
- In a second step the user can search for Sensor Observation Services in the catalogue supporting the phenomena he is interested in.
- In the final step the user can directly access the URLs of the Sensor Observation Services provided in the catalogue results to access the observations contained in the SOSs.

The principle is illustrated here:



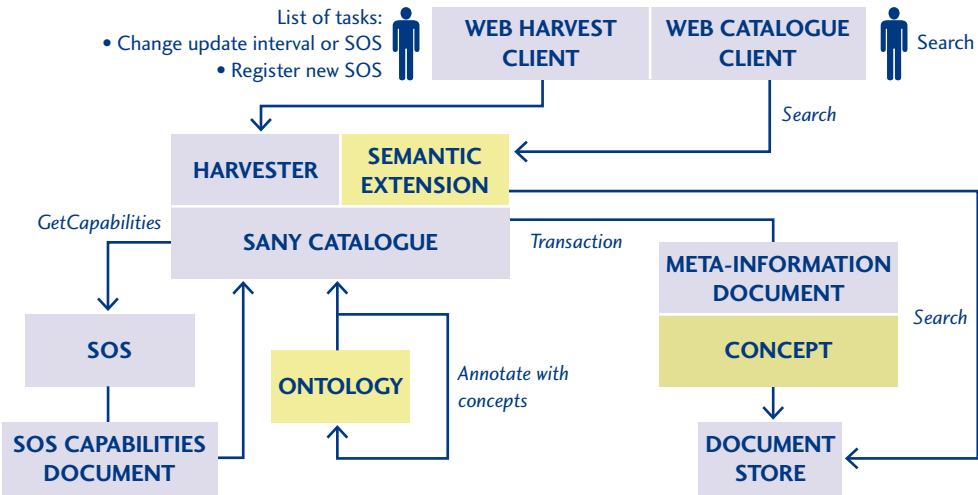
Another research topic for the SANY catalogue was the automatic creation of meta-information. Since the meta-information schema was designed according to O&M, which is also used by the SOS, it is possible to use its operations GetCapabilities and DescribeSensor to automatically harvest the meta-information, which is

necessary to create SANY meta-information documents. The catalogue service has been extended with a harvesting operation for this task.

Besides the automatic creation of SANY meta-information documents, the automatic creation of INSPIRE meta-information is also possible. In this case the information provided by the SOS is not sufficient for the creation of instance documents compliant with INSPIRE schemas and must be extended. For more information on this, please refer to chapter 8.1.3.

To overcome problems with the discovery of unharmonized URNs used in phenomenon or feature of interest descriptions of an SOS, the principle of semantic annotation has been tested. In the test example a SOS provides links to an ontology and a lifting schema, which describes the relation between the ontology concepts and the SOS phenomena. In order to provide flexibility, the W3C recommendation ‘Semantic Annotation for WSDL and XML Schema’ (SAWSDL) has been used (Farell, Lausen). The harvesting operation infers from the phenomenon to the related ontology concept and includes the concept into the created meta-information document. The catalogue client provides the user access to the used ontology. The advantage is, that a user using the ontology concepts for his search will get more results than a user performing a search with a-priori knowledge of phenomena available in the catalogue: a search for the observable property ‘relativeHumidity’ leads to results of a specific SOS. A search for ‘rf’ leads to results of another SOS using this observable identifier for the very same phenomenon. But a search using the ontology concept ‘relative_moisture’ related to meteorology leads to results of both SOS.


The illustration below shows the harvesting architecture of the SANY Catalogue Service in combination with the semantic annotation:



MAP AND DIAGRAM SERVICE

7.8

The Map and Diagram Service is an example of a cartographic Web service. It can be defined as a service that visualizes, symbolizes and enables the geographic clients to interactively visualize topographic and thematic data. The main task is to transform geographic data and thematic data, including geo-referenced sensor data, into a graphical representation using cartographic rules. The illustration below shows an example of sensor data visualisation with a colour map and contours:



The Service Oriented Architecture allows seamless information integration by abstracting the complexity of the heterogeneous nature of the data sources. In this context, sensor data, as handled within the SANY Project, serve as a good illustration for the (dynamic) nature of spatial information that must be represented in the form of maps. Modern cartographic applications are required to immediately reflect the updates in the data without sacrificing the cartographic quality. This novel situation has a considerable influence on the established cartographic workflow. In order to produce the map, cartographers do not have the possibility anymore to prepare and symbolize directly the data. They might even know that certain parts of the data are being changed or updated on a continuous basis, as in the case of sensor data. Therefore data

symbolization has to be thoroughly controlled in an open and distributed manner by Cartographic Web Services.

The main design consideration is not to replace existing standards, but to extend them for cartographic usage. To support interoperability and use of open standards, the Map and Diagram Service is based on and enhances OGC standards. The Map and Diagram Service Specifications introduces several operations based on the Web Map Server (WMS), Styled Layer Descriptor (SLD) and Symbology Encoding (SE) standards. The well known WMS requests can be recognized in following presentation of the Map Diagram Interface operations:

- `getMap` returns a map of spatially referenced geographic and thematic information as an image document.
- `getDiagram` returns a diagram representation of tabular data as an image document.
- `getFeatureInfo` returns information about the features rendered in a certain point of a map or diagram layer.
- `getLegendGraphic` returns a legend symbol corresponding to a layer as an image document.
- `getLayerDescription` returns a layer description document containing schema information for a layer.
- `getStyle` returns the cartographic rules (style) associated with a layer.

The Map and Diagram Service interface specification follows and complies with the WMS 1.3 specifications and the SLD profile for the WMS. Symbology Encoding and the Styled Layer Descriptor Profile for the Web Map Service Implementation Specification are the direct follow-up of the original Styled Layer Descriptor Implementation. SE is the most recent OGC standard for portrayal of geographic information.

The combination of Web Map Services, Styled Layer Descriptor and Symbology Encoding already provides a viable open framework for basic topographic representations. However, advanced cartographic features like user-defined point symbols, multi-layered symbols, transparencies, textures, marking, patterns and diagrams are the means that enable the cartographer to achieve map quality required by environmental management applications: a good differentiation of features and map legibility. In this direction OGC WMS and SLD standards are generally considered as too restrictive. Absence of custom vector-based point symbols, patterns for spatial features and layer transparencies limit the usability of WMS from a cartographic perspective. Moreover, its inappropriateness to create thematic maps is the main reasons why WMS is used for presenting topographic maps, and not for thematic representations.

Fortunately, these standards are very flexible and can be cartographically enriched to fulfil the complex visualisation requirements coming from environmental management. As such, the Map and Diagram Service implements several extensions (documented in the OGC Change Request 07-105) as a solution for cartographic challenges of environmental management.

WEB PROCESSING SERVICE

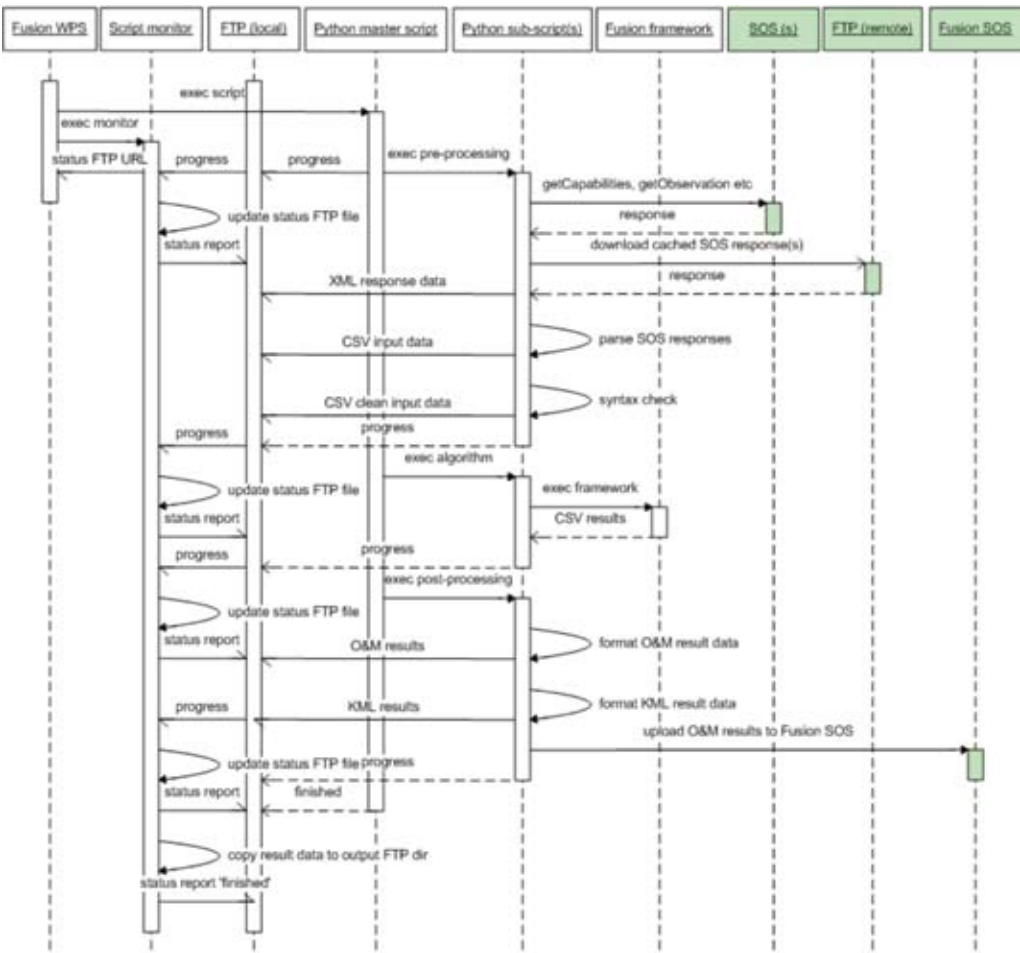
7.9

The Web Processing Service (WPS) is an OGC standard interface for a processing service and has been used to provide fusion and modelling services within SANY. In SANY we use WPS v1.0.0, implementing support for both complex and literal WPS inputs, and both reference and literal WPS outputs. Status reports are returned via a FTP site and the final result sets are returned via both a FTP site and a Fusion SOS.

The following operations are specified within the WPS standard:

- `GetCapabilities` for requesting a self-description of the service.
- `DescribeProcess` for requesting the list of processes supported by a specific WPS instance; this includes information on the input parameters and expected output results.
- `Execute` for initiating a new processing action.

There are two WPS implementation case studies in SANY. Both implement v1.0.0 of the standard and provide results either literally or via O&M formatted result sets. The underlying Execute operation behaviour is implementation specific, enforcing only that the WPS protocol is observed of providing either an immediate literal result or an URL to a XML status file for client polling of progress.



By way of example one SANY implementation of the Execute operation spawns a fusion Python script process to run each new fusion algorithm. In addition, a handler process is spawned to monitor data fusion progress and update the status XML file located on an FTP site. Once finished, the data fusion output is made available on the FTP site. The URL link to the FTP is written to the final status report which the client receives. Results are formatted in both low level comma seperated value (CSV) format and SWE O&M format for upload to a Fusion SOS.

TIME SERIES TOOLBOX

7.10

Sensors often record one specific observation repeatedly over time, and applications in sensor networks have to store and process such data, which is also called 'time series'. The simplest form of a time series is a single floating point number, e.g. temperature recorded at regular intervals.

The Time Series Toolbox (TS Toolbox) is a set of software components and application programming interfaces that simplify the task of building applications that record, process, store and publish time series of observations. The TS Toolbox contains software components for the following functional areas:

- Data connector components implementing access to data using various protocols and data formats
- Core components interfacing with the connector components and providing specific additional functionalities like data processing or caching
- Frontend components implementing interface functionality (user interfaces or software interfaces)

The functionalities implemented by TS Toolbox components provide application developers with higher-level building blocks than typical general purpose libraries, and allow rapid development of fully fledged applications. The TS Toolbox also includes example applications that can be either used as they are, or as a basis for developing more complex applications. The following components are included in the TS Toolbox:

SOS FRONTEND	DATA PUMP	GTV	...
FORMULA3	TIMESEIES API		CACHING
SOS CONNECTOR	CSV CONNECTOR	ANYSEN CONNECTOR	UWEDAT CONNECTOR

The ‘frontend’ TS Toolbox components provide interfaces to users or other applications. Currently, three frontend components exist:

- The SOS frontend component simplifies the task of developing applications with an OGC Sensor Observation Service compliant interface. It provides an implementation of a SOS service on top of the Time Series API, and is based on the 52° North SOS implementation.
- The Data Pump frontend component implements functionality for transporting time series data from one data connector component to another. As such it can be easily used for creating applications to import and export data between applications.
- The GTV frontend component implements functionality for building GUI client applications for accessing and displaying time series data.

The ‘data connector’ TS Toolbox components provide access to data using various protocols and data formats; this includes three general purpose data connector implementations:

- SOS connector: by using this connector implementation, applications can be interfaced to OGC SOS compliant services. Currently the SOS connector supports reading data from a SOS service, and storing data in a SOS service using the transactional profile of the SOS specification.
- CSV connector: many legacy applications implement functionality to export data in simple comma- or tab-separated text files. The CSV connector component allows seamless integration of this type of data in TS Toolbox-based applications.
- AnySen connector: an implementation of a data connector fetching data from a sensor driver that interfaces with physical sensors. The AnySen connector implements a flexible configuration scheme allowing it to be adapted to different vendor protocols.

The TS Toolbox also includes one example of a ‘legacy connector’, which is used to access the air quality data stored in the proprietary data acquisition system ‘UWEDAT’ monitoring systems. Legacy connectors allow much tighter integration of legacy applications, e.g. real time access; storing altered data back to original service etc., than exporting and/or importing data to integrate those applications.

The TS Toolbox currently includes two reusable implementations of ‘core’ components:

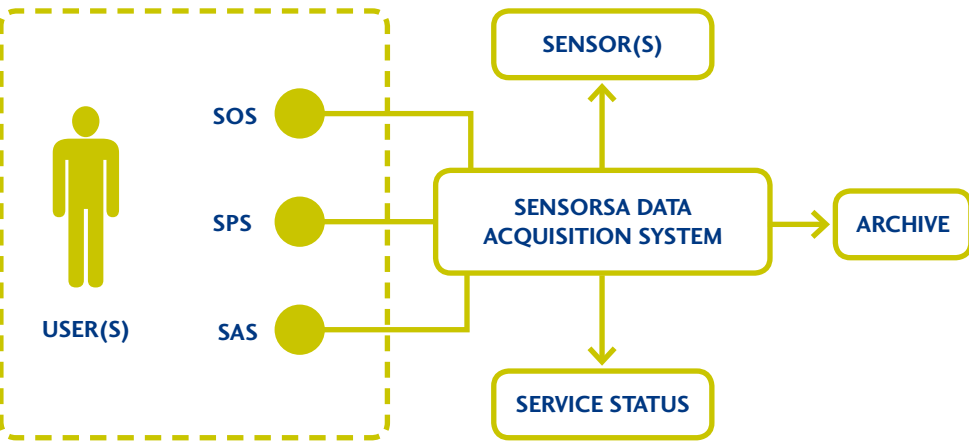
- Formula 3, a concise, text-oriented, high-level language for manipulation and transformation of time series data, enabling users to efficiently implement processing logic.
- A caching component, which allows temporary storage of the data within an application and offers an easy way for including pre-fetching and caching capabilities in applications and services.

In SANY the TS Toolbox components are used in the following applications: the Cascading SOS Service (section 7.2), the SensorSA Data Acquisition System (Section 7.11), GTV (section 7.12), and in the Universal Data Pump – a simple application that provides a convenient way for transporting data from one application, service or file for which a TS-Toolbox data connector exists to another.

SENSORSA DATA ACQUISITION SYSTEM

7.11

The SensorSA Data Acquisition System (SensorSA DAS) is a network capable appliance developed by AIT, which allows seamless integration of various sensors in a Sensor Service environment based on SensorSA. The main characteristic of SensorSA is the exclusive use of OGC SWE interfaces for all communication. SensorSA DAS exposes sensor data, management data, and history of alerts over a Sensor Observation Service (SOS) interface. The sensor configuration is performed via Sensor Planning Service (SPS) interface, and the configuration of events and alerts is performed through a Sensor Alert Service (SAS) interface:




The SensorSA DAS aims to support the ‘plug and measure’ type of operation foreseen by SensorSA at the sensor level. This means, that a sensor plugged into a SensorSA DAS should be immediately recognized, configured and integrated in a SensorSA network. This can only be done for ‘smart sensors’, i.e. for devices that provide some kind of an ‘Electronic data or product sheet’, which can be automatically read and interpreted by the DAS.

Ideally, the sensor provides all information required to configure SensorSA DAS using an electronically processable sheet. Since typically most sensors only provide a small subset of the information required for configuring the SensorSA DAS, within SANY *the description offered by sensor is used as unique key for finding the corresponding configuration data in a database of ‘known sensor types’.*

7.11.1 SensorSA Smart Sensor Adapter

The SensorSA ‘Smart Sensor Adapter (SSA)’ is a simple device which allows the user to connect a simple RS-232-based measuring device with automatic identification and registration to the station computer.



The SensorSA SSA has a possibility to provide all information required for automatic configuring of a SensorSA DAS, including e.g. capabilities of the sensor, resolution, accuracy and type of the measurements (units), sensor location, owner, proposals for information processing and more.



When attached to an USB port, the SSA allows the DAS to download its configuration data. In a next step, the DAS uses this information in the similar way it would use the configuration data obtained from the ‘sensor types’ database. Finally, the SSA switches to the ‘transparent’ mode and allows direct communication between the SensorSA DAS and the SSA RS-232 interface.

The current implementation of the SSA is based on a Microchip evaluation board shown above and the firmware supports following operations:

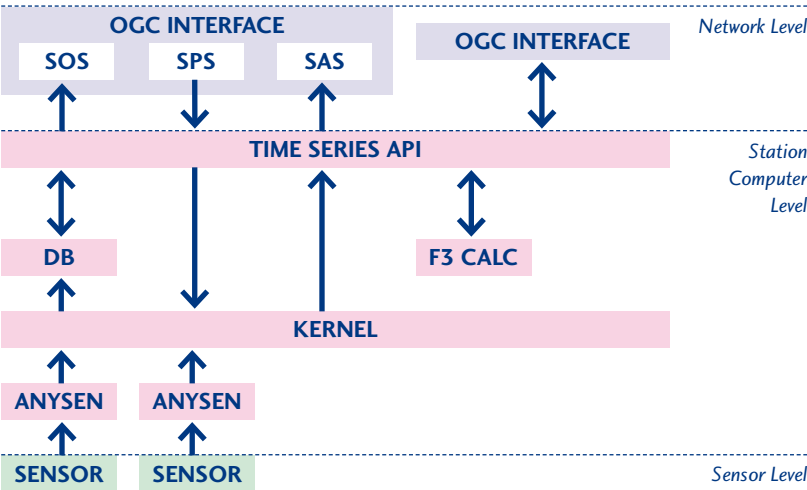
- **setTransparent=0 or 1** switch to Command mode or RS232-USB converter mode
- **getCapabilities** outputs the Sensor-ML file to USB
- **getTemperature** outputs the local temperature of the adapter (demo-value)
- **getSwitchStates** outputs the logical states of two digital input lines
- **getPoti** outputs the value of the on-board potentiometer (demo-value)

These commands can be sent over the USB connection and will be interpreted by the board until it enters the transparent mode (setTransparent=1 command)..

Once in transparent mode, the SSA acts as a simple USB to RS-232 bridge, which allows re-using of the existing sensor drivers with no or minimal changes.

AnySen Driver

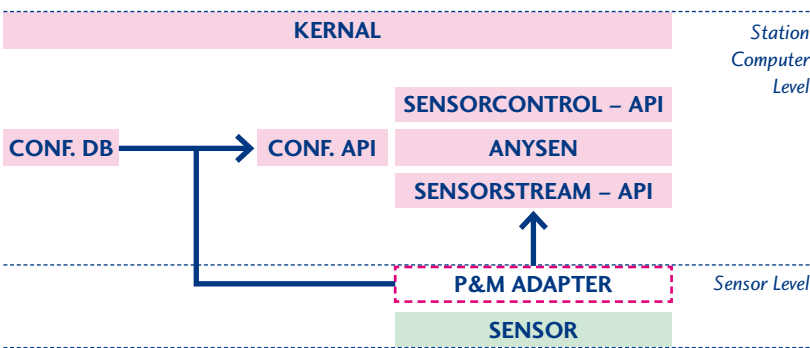
The AnySen driver is a software component for low level data acquisition, which can be used to connect sensors to a Data Acquisition System. It controls the sensor, interprets measurement streams and parses measurements. The following illustration shows the place of this component in a typical DAS:



AnySen is a part of the Time Series Toolbox and can be easily combined with other toolbox components. The main advantage of AnySen is its capability to read and interpret data from many sensor nodes equipped with a digital interface (e.g. RS232 or LAN). This is achieved by abstracting the sensor protocols and reading the concrete description out of a simple sensor description file. All commands necessary for configuring and retrieving data from an analyser can be configured at run time.

When a new sensor is attached to the DAS, the AnySen can trigger an automated configuration process, leading to the sensor level 'Plug and Measure'. The configuration can either be read from a central Database, or from the SensorSA Smart Sensor. The configuration possibilities include the protocol description, structure of the measurements, and various meta-information such as the name and unit of the measured phenomenon. This is a significant difference to current DAS concepts, where this meta-information is injected in higher levels, often as part of the application logic.

The following illustration shows the concept of AnySen in detail:



The **SensorControl API** provides an interface for controlling sensors and retrieving single measurements. This allows easy integration of additional sensor-drivers to the kernel, e.g. for drivers specialized at complex sensors producing spectres, 2D- or 3D coverages.

The **SensorStream API** provides the technology-independent interface for communicating with the sensors and assures the basic protocol handling is separated from the central AnySen logic. Current version of the AnySen driver has been designed for sensors connected over RS232 interface, but the SensorStream API allows easy integration of the sensors communicating over other interfaces as well (e.g. TCP-IP, ZigBee, or CAN-Open).

The **Configuration API** provides the technology independent interface for AnySen configuration. This allows easy adaptation of the AnySen driver to

various configuration file formats (e.g. SensorML, LUA, YAML, ...). In addition, the Configuration API also provides a simple mechanism for mapping of the semantically equivalent configuration options (language-, domain- or vendor specific naming conventions). For example, the AnySen requires a configuration parameter specifying the length of the measurement string. Internally, this property is called 'ms.length', but this property may be called 'Messwert.Laenge' in the configuration file.

Conf. DB is a database with configuration data for all kinds of sensors. A local copy of this database resides on each DAS. The required configuration files can be downloaded from a central database on request. The network-wide consistency of the configuration can be assured by an update-mechanism.

The **Plug and Measure Adapter** is an optional add-on-device, which is permanently attached to an analyser to ensure its automatic detection and configuration by the host system. One example of such a device is the SensorSA Smart Sensor Adapter.

GENERIC TIME VIEWER

7.12

The Generic Time Viewer (GTV) is a generic desktop application and a toolbox for building specialized applications capable of presenting a common and combined view on time series data stemming from different sources, such as sensors, simulation models or data fusion outputs.

The GTV is implemented in Java on a rich client platform. It is an expert tool for the daily work of decision makers mainly in environmental authorities. The development of the GTV has been started within the SANY project, and the design strongly reflects the requirements inherent to the air quality monitoring domain. Nevertheless, the GTV can be easily adopted to the needs of other environmental domains. The main design goals of GTV were:

- to develop an expert tool capable of accessing and visualization of all data used within the SANY project;
- to assure the GTV provides efficient and reliable support for domain experts inspecting large amounts of data.
- to assure the GTV is easily extendible in order to answer the future user requests for additional functionality

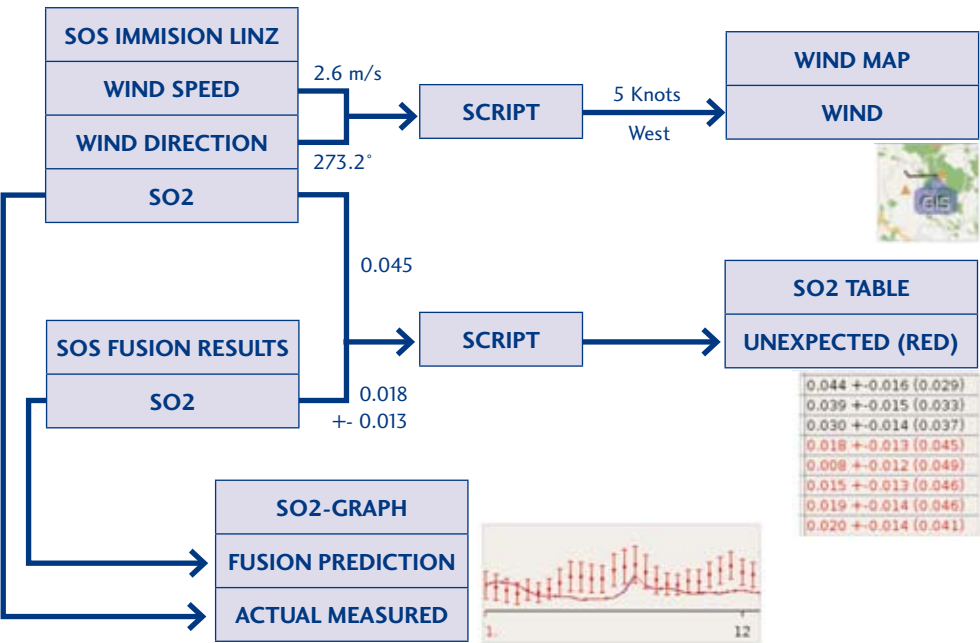
The main GTV components are mainly a set of connectors to remote systems and a set of viewer windows, which can be combined and configured in a flexible way. Both, connectors and viewers can be easily added at runtime without the need to recompile or reconfigure the application.



As illustrated above, the GTV currently provides three viewer components for visualization of the data in tabular form, as x-t graph or on a map, but new connectors and viewers can be easily added at runtime without the need to recompile or reconfigure the application.

The most interesting GTV feature is the possibility to process the data on the fly using the build in scripting features. In addition to evaluating the data from one or more sources, the script decides which visualization component(s) are invoked to display the result. For example, the data from two sources can be compared and the differences visualized using the symbols on a map or colour-coded tables.

Depending on the acutal configuration, the views may be either independent, or connected and capable of dynamically synchronising their context. For example, the symbols on a map (wind speed and direction) can reflect the time chosen by the user in one of the other two graphs, and browsing through data in tabular form can result in animated maps. The figure below illustrates the GTV user interface with three views of the data used in Air Quality pilot (Section 8.1).




The GTV is currently used as the basis for the 'Data inspection client' in Air Quality Monitoring (SP4) Pilot (see section 8.1). In the future, the existing processing component may be replaced by more powerful TS Toolbox Formula 3 processor component, and the performance improved through inclusion of the caching component.

IGN GEOCUBES

7.13

Within the context of SANY, IGN has implemented an innovative risk management application in the field of geo hazards. The team who works on this project has decided to focus on the detection of subsidence or landslide, using a network of mini-GPS devices, also called 'Geomotes'. These GPS are extremely small (about 30 x 40 x 10mm, antenna included) and can provide an accuracy better than one centimeter, thanks to a post data processing based on a differential calculus. These devices will be embedded in self-powered systems which are connected wirelessly to each other in order to set up a network. Each node, named Geocube, can continuously provide GPS data to one PC. Furthermore, Geocube is equipped with a three axis MEMS accelerometer which is able to detect high frequency displacements. Therefore Geocube is able to send a warning, then GPS data are being recorded and new positions are computed with one hour delay.



A network of Geocubes is a mesh with wireless links that can monitor a local area of about one square kilometer. This network is small, easy to install and a low cost solution for geohazard predictions. Each node, i.e. each Geocube, is geo-localized with a relative positioning accuracy better than one centimeter in planimetry and two centimeters in altimetry. Moreover, time accuracy of each Geocube is better than 50ns: it can easily and precisely date or initiate events.

7.13.1 Description of a Geocube

A Geocube, as its name suggests, is a cube shaped device, which supports a set of solar cells and contains three Printed Circuit Boards (PCBs) and a battery pack. Below each Geocube, three external connectors allow the user to link the Geocube to different types of external devices.

Since the main role of the Geocube is to allow the transmission of GPS data, its architecture is made around the mini-GPS and the communication module. Since knowledge about ambient air pressure can be useful for the data processing, the board also contains a pressure sensor. In order to limit the power consumption and extend the battery life, the GPS functioning with it's comparatively high power drain is not used permanently. Instead an accelerometer has been added, which runs permanently and switches on the GPS module to sense motion in case of any abrupt movements. The Geocube interfaces with a computer via USB and RS232.

The battery pack can be charged either by the solar cells or by an external source. In addition, components are integrated to protect the battery and provide access to its charging status.

A Geocube has three small (42mm x 44mm) solar panels on three sides of the cube – the fourth side is purposely positioned facing a north direction. Each solar panel provides up to 40mA under 6V depending on the solar illumination.

If required, a bigger external solar panel can be fitted, e.g. for Geocubes that are positioned in the shade. The dockable connector under the Geocube allows plugging external power source with different voltage: 3.3V, 5V and 4.5-28V.

7.13.2 Provided Data

A standard Geocube without any peripheral or external sensor can provide these five data using wireless communication or via USB or RS232 cable:

- GPS data: RAW data measurements: data provided by the GPS, contains carrier phase, pseudo range and Doppler measurements (40 to 350 bytes)
- Pressure (2 bytes)
- Temperature (1 byte)
- Three dimensional acceleration (2 bytes/axis)
- Battery state of charge (2 bytes)

7.13.3 The Network of Geocubes

Among the wireless communication protocols, Zigbee appears to be the most suitable for the sensing devices deployed in the SANY project. Indeed, it is intended for use in embedded applications, which require low data rates and low power consumption. It is also simple to design, reliable and interoperable. Nevertheless, Zigbee is not really adapted for mesh networks because of the impossibility for routers to switch to stand-by. That's why some manufacturers have developed other very similar protocols with stand-by router capability.

A network of Geocubes is composed of one Geocube set as coordinator. This one is directly interfaced with the computer that collects the data. Then, depending on their location, other Geocubes will be either set as router or end device:

- The end device sends its data every 30s during one or two hours a day and switches off when it is not transmitting, i.e. most of the time. If an abrupt movement is detected by the accelerometer, it warns the coordinator and sends its data every 30s until the coordinator asks it to stop.
- Router: as well as doing the same actions as those of the end devices, a router passes data from other devices.

All the devices of a network periodically switch on at the same time in order to re-synchronize and to transmit potential alert messages.

If the network is jammed for any reason, each Geocube can save its data in an embedded microSD of 1GByte.

This small flash memory card is very useful to upgrade the firmware of each Geocube. Indeed it is possible to send a new firmware using the wireless communication to any Geocube of a network; this firmware is written on the microSD and as soon as this Geocube restarts, it compares the version of the installed firmware with the downloaded firmware and installs the new one.

7.13.4 External Devices

Each Geocube has two coaxial connectors to emit a precise time pulse or to date an external event using the GPS module.

The third connector is a dockable connector with 30 pins. Thus it is possible to link a Geocube with another external device with different sensors like vibration sensors, sonometer, bathymeter...

7.13.5 Post-Processing of GPS Data

Most civilian GPS chips which are integrated in car navigation systems are single frequency receivers. They generally use only the Coarse/Acquisition (C/A) code, which prevents them from providing a positioning accuracy of less than one meter.

The GPS chip contained into each Geocube is also a single frequency GPS receiver that can correlate on C/A-code but also measure carrier phase on L1.

A high positioning accuracy can be achieved by working in differential mode, where additional positioning information is provided by at least one reference dual frequency GPS station.

Both, C/A and carrier phase data with Doppler measurements can be easily post-processed in differential mode. The setup of a dual frequency reference station supports the use of ionospheric corrections in the calculations. Within SANY, IGN has developed a software to process GPS data and detect movement of one Geocube with respect to the mesh network.

7.13.6 Applications

The main application using Geocubes is landslide monitoring. In this case, movements are very small over a very long time before the major slide sets in, so these displacements have to be continuously monitored. The installation of a network of Geocubes is very easy and can be very quickly deployed on risk areas. The network is totally autonomous and each Geocube can be upgraded through wireless communication.

Other applications have also been considered. These include marine swell measurements by setting Geocubes on buoys. It is possible to measure the amplitude and frequency of wave swells, e.g. in order to find the best site to install a tidal power station.

Geocubes may also be very interesting for oil-prospecting and potentially replace the current geophones, which are still linked by very long cables.

SOLDATA MICRONS

7.14

In order to provide reliable but easily and quickly deployable sensor nodes for geotechnical applications, SolData has developed a smart sensor system within SANY, called 'Microns'.

'Traditional' geotechnical monitoring systems are mostly wired systems, with a pre-defined number of sensors organised according to an instrumentation plan. The installation of such systems, whether considering total stations or borehole sensor systems installations is often time consuming, and the wired networks are not flexible in their configuration. In consequence, such networks cannot be easily and quickly extended in the case of crisis or higher risks.

In environmental risk applications, such as landslide monitoring, the areas of interest are often difficult and inaccessible, far away from communication infrastructures, and with no permanent power source. In such conditions, the installation of a 'traditional' geotechnical system, if not impossible, is problematic and induces logistics efforts, which can be costly. Also, access to the monitoring zone may pose issues of health and safety for staff. In such cases, wireless sensor nodes can be quickly deployed on site and essentially operate autonomously over a long period of time.

After a survey of wireless sensor products on the market, none of the available sensor nodes were really suitable for geotechnical applications. Most devices do not permit the use of a wide variety of sensors and field applications since they were designed for specific industrial applications, with no resilience to harsh construction environments and extended outdoor exposure.

As a result, SolData took the opportunity of developing wireless sensor communication nodes, organised in a self-configurable and autonomous network for a dynamic and adaptable sensor system.

Description of a Micron

7.14.1

As opposed to the Geocube, the Micron is a sensor node and not an actual sensor. Up to eight sensors can be connected to a Micron, which in its turn stores and relays observations to the sensor network.



The Micron smart sensor nodes have an integrated battery, which allows a guaranteed autonomy of 2 years. This autonomy may vary depending of the sensors attached and the frequency of measurements. An external power source can be connected, so e.g. photovoltaic panels could also be deployed to power the sensor nodes, and take over the battery powering.

The voltage supported by the devices are :

- +3.3V/25mA,
- +5V/25mA,
- +9V/25 mA,
- +12V/25 mA,
- +15V/25mA.

7.14.2 Provided Data

The data provided by the Micron smart sensor nodes depends on the type of sensors that are connected to it. A Micron supports between 4 and 8 sensors of different types, depending on their type (polar or differential). Therefore a smart sensor node is not dedicated to a specific type of sensors, but can be used as a central acquisition point for several kinds of sensors, offering a better flexibility to the users' needs in term of monitoring sensor system.

Each sensor node has 4MBytes of integrated memory, which can store up to 184.636 acquisition messages.

Network of Microns

7.14.3

As the sensors need to be rapidly deployable in any site configuration, a wireless protocol must be used for flexibility and to avoid problems due to wiring and power source limitations.

The SolData smart sensor system is based on wireless communication, supporting the ZigBee wireless communication standard with several features including multi-hop, self-configuring, self-healing and dynamic routing. This protocol, requiring very few instructions, facilitates a low power consumption and high autonomy, the integration of a great number of nodes, and has an acceptable communication range for our application scenarios. It operates at a radio frequency of 2.4 GHz, therefore no radio licensing is needed to make use of those devices.

The sensor nodes support different network topologies and are organised in a self-healing network, which means that if a node fails, the other nodes will automatically find another communication route. When several routing paths are possible, they will chose the most efficient path for their data communication.

The communication range of the existing prototypes is 300m (line of sight), but with the latest developments this range should be greatly improved.

TRISKEL MARINE LTD. DATABUOY

7.15

Triskel Marine Ltd is a UK based marine data management company, specialising in gathering a wide range of data from the marine environment and transmitting it ashore for processing and analysis. SANY used TML's marine data monitoring buoys to demonstrate the ability to collect data from mobile marine sensors and to combine this data with information from other sources. The combined data was used to produce decision support tools for two rivers in Cornwall, UK.

The project collected real time water quality and current data from the Fowey and Fal estuaries over several months. This was then combined with meteorological data from the river basins and historical data from the shell fisheries. The merged data set was used to produce a prototype mathematical model for predicting microbial bloom in the rivers.



Triskel Marine’s DataBuoy is a cost effective way of monitoring the inshore marine environment in real time. The buoy is tough, light weight, easily deployed, and completely autonomous, with the following specifications:

- Diameter/height: 610mm/305mm
- Total mass (ballasted): ~30kg
- Draft (ballasted): <1 m
- Measurement depth: To user’s specification
- Interfaces: 1 serial, 3 analogue and 1 digital
- 7Ah battery and two 10W solar panels
- Global Positioning System (GPS)
- (GPRS) modem for bi-directional telemetry, reporting at a maximum frequency of 12 times per hour. Text messaging and email are available as standard

7.15.1 Provided Data

The data boy is capable of accepting inputs from a wide range of different sensors, observing parameters such as turbidity, salinity, temperature, dissolved oxygen, pH, current speed and direction, depth and meteorological parameters. Onboard sensors include:

- YSI 600 OMS V2 optical monitoring sensors measuring
 - ◆ Wiped optical turbidity sensor, 0 to 1000 NTU +/- 2%
 - ◆ Water temperature -5 to +50 deg C +/- 0.15 deg C
 - ◆ Salinity 0 to 70 ppt +/- 1%
- Airmar ultrasonic current sensor
 - ◆ Speed range 0.1 – 40 knots at 2 Hz update frequency
 - ◆ Transmission frequency 4.5MHz
- Autonnic Research floating core magnetometer

Network of Data Boys

Real time environmental data is transmitted using GPRS technology at intervals from 10 minutes to 24 hours. The standard unit needs no setting up, just placing it into the water. It is ideal for collecting long term trend data and for monitoring transient events such as dredging, spills and other pollution. Communication with the buoy is two way – updating factors such as frequency, alarm levels and calibration constants, can be set via a website.

Alerts, triggered when an alarm threshold is exceeded, can be provided by text message and email.

There is no limit to the number of buoys that can be deployed simultaneously and each is individually addressable from the website of the user. All buoys are fitted with GPRS, GPS, battery monitoring and solar panels as standard. Weather reporting, on-board data logging and navigation lights are available as an option.

7.15.2

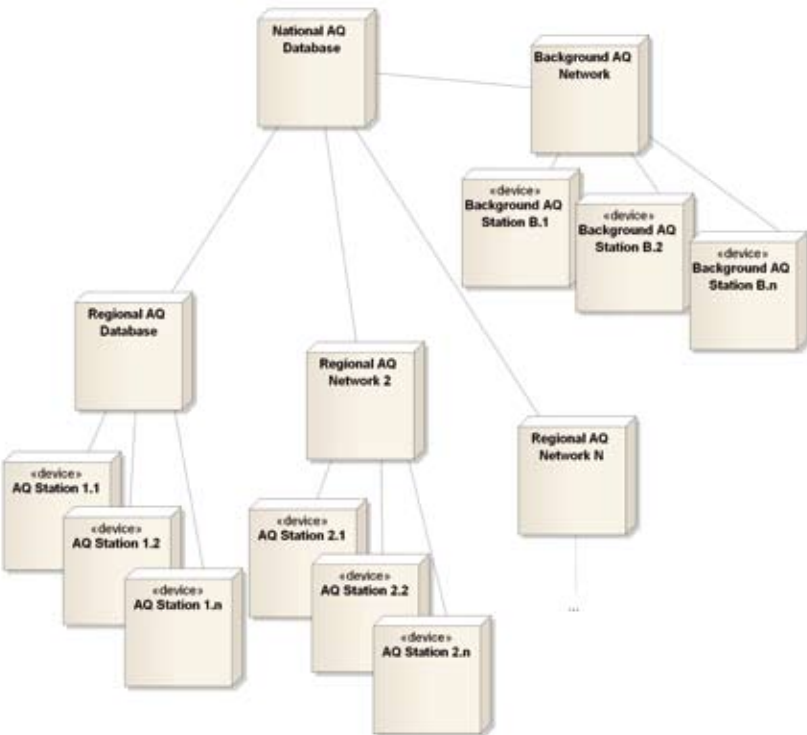
8 SANY Applications

This chapter is written for those who want to see how SANY components and the SensorSA can be used in real world applications. It will address the following questions from an application point of view:

- What are the Pilots about?
- How were they implemented?
- Which services and SANY components were used?
- What are the benefits of doing it the SANY way?

8.1 AIR QUALITY MANAGEMENT

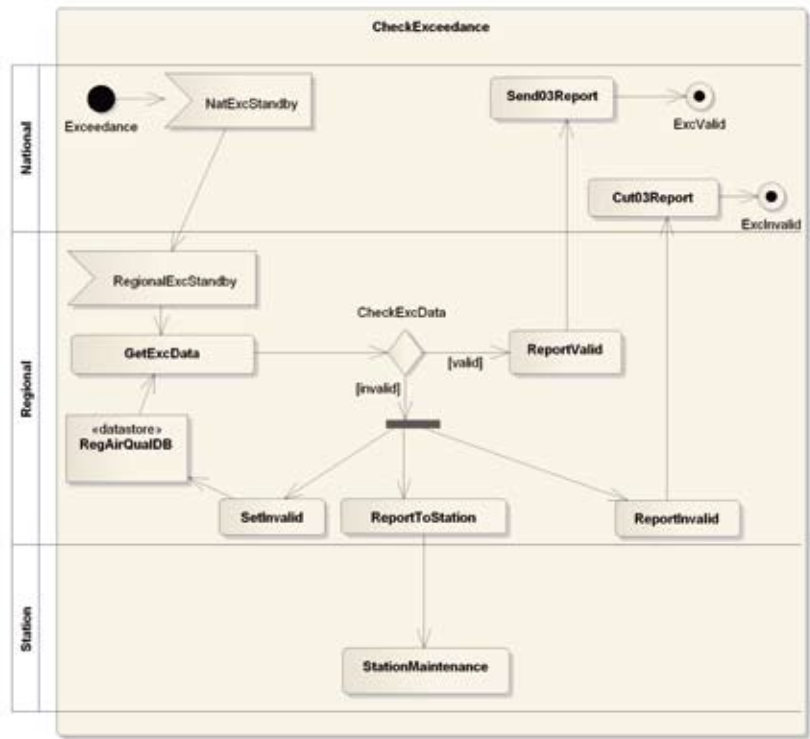
The Austrian air quality network is organized as a decentralized system, mirroring the federal structure of administration in Austria:



Each of the nine regional provincial governments operate a regional air quality monitoring network within their federal province. In addition, the Umweltbundesamt, the expert authority of the federal government in Austria for environmental protection and environmental control, also maintains a network of ‘background’ measurement stations. These stations are positioned in natural habitat far away from the main roads, industry and settlements.

Each network owner is responsible for the entire quality assurance and storage of measurement data, using reference standards for calibration, which are provided by the Umweltbundesamt. The resulting data is then transmitted from the provinces to the national air quality data-base operated by Umweltbundesamt, and to the European ‘Near Real Time Information System’ operated by the European Environmental Agency. The data from this process is used for generating provincial and national reports, including those submitted to the European Commission.

Austrian Air Quality Monitoring data is subject to multiple quality controls at different levels of administration. The basic QA procedure involves manual inspection of the data by domain experts, and marking the data as ‘valid’ or ‘invalid’. The figure below illustrates a simple quality control process, used to check the validity of exceedance.



When an exceedance is detected, the national air quality expert on standby is automatically notified via e-mail and SMS. The national air quality expert in turn notifies the regional expert that an exceedance has occurred, and requests input from the regional expert.

If the regional expert declares this exceedance of threshold to be 'valid', this information is returned to the national expert, who in turn gives the clearance for transmission of the values for reporting purposes.

If the limit exceedance is deemed as invalid, this value is flagged as invalid within the regional AQ DB (air quality data base) and the national expert is notified. In addition, the station operator is informed of the invalid values being generated by the station. Finally, the updated data marked 'invalid' is uploaded to the national air quality database.

8.1.1 SANY 'Air Quality Management' Pilot

The SANY 'Air Quality Management' pilot illustrates how SANY results can be deployed in the context of Air Quality Management. One of the objectives in this use case scenario is to extend existing systems with state of the art components without the need to replace the pre-existing sensor network infrastructure.

One pilot implementation is located in the vicinity of the City of Linz, Austria. This pilot site features two main types of sensor data:

- immission measurements at 17 locations, and
- emission measurements from major industrial plants in and around the City of Linz.

The existing air quality monitoring system is based on UWEDAT, an environmental monitoring system based on Windows NT/2000 and mainly used in the field of ambient air quality monitoring.

The system provides access to real time measurements as well as several years of archived data. Additional immission data can be gathered with SensorSA Data Acquisition System (DAS) prototype, or imported from external sources.

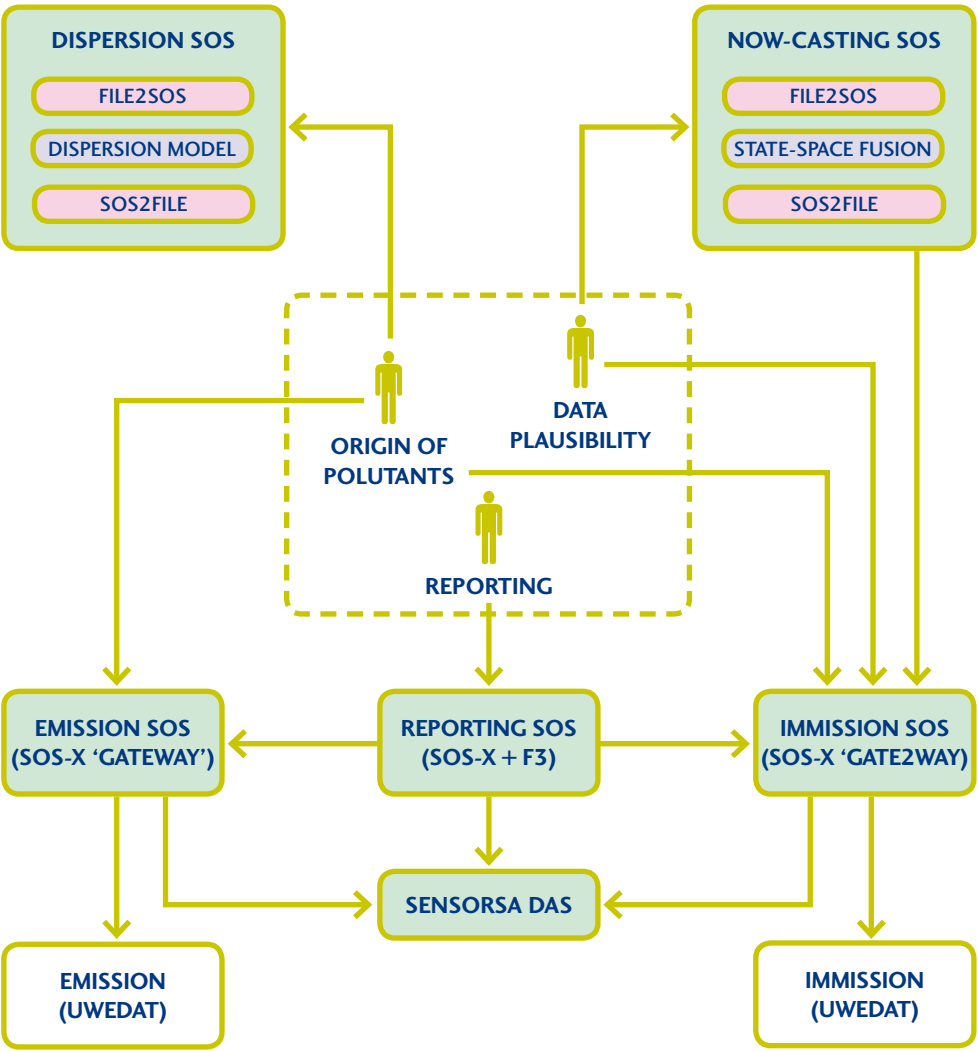
For the SANY Pilot, the concept of the Cascading SOS has been implemented to add a standards based interface to the existing UWEDAT system.

The data is used by three other SOS instances as part of the SOS-cascade that provide additional processing:

- 'now-casting SOS',
- 'Dispersion SOS', and
- 'Reporting SOS', based on the cascading SOS concept.

The figure below illustrates the use of SensorSA DAS, Cascading SOS, fusion- and modelling- services in Linz pilot. For simplicity, the figure omits all other SensorSA services and end-user applications. Thick arrows connect data consumers with their main data sources. Alternative and optional paths are represented with thin arrows:

The information provided by many of these services can be either visualized using an SOS compliant client, or used as the input for further data-processing services and special purpose clients.



8.1.2 Cross-border Data Integration

Existing air quality monitoring systems are often implemented as proprietary networks designed for a particular purpose, following national or regional specifications. This, of course, limits the capability to directly interact with other systems for purposes, not anticipated in the design phase. Real time cross-border usage of measurements to support pan-European environmental management is a typical example. It is also a typical challenge on the IT side, since existing systems with similar purpose but most likely different implementation approaches need to be opened up to become interoperable on a higher level. At the same time there should be no significant interference with the existing operational systems.

To address these issues, SANY has implemented a second Pilot in Flanders, covering a cross-border region between Belgium and France. The objective was to investigate the feasibility of cross-border data exchange from existing unrelated systems using SensorSA infrastructure. It possesses the following characteristics:

- It involves 2 physically distinct Monitoring Systems.
- The Monitoring Systems are operated by 2 separate Monitoring networks located in 2 different countries (France & Belgium)
- Both systems use similar technology, but due to administrative reasons cannot share data in real time.

So even with similar underlying technologies, a combined direct access in real-time to existing data is impossible, in this case because of administrative issues. To build a bridge, each of the existing systems has been wrapped with an OGC compliant Sensor Observation Service in order to achieve interoperability on a cross-border basis. The inherent capability of the SensorSA is then used to assure that the data can be easily found and accessed by all relevant users, independently of its origin.

The main advantage of the SensorSA architecture in this respect lies in the possibility of cost-effective reuse of existing sensor infrastructure. There is no need for the implementation of a new air quality monitoring system from scratch or the development of static proprietary ad-hoc bridges between each system. The use of the standardized SOS interfaces on top of the existing monitoring system provides dynamic access to the respective data, as illustrated in the screenshot of a client application below:



INSPIRE Meta-Information

The 'Cross Border Data Integration' example described in the previous section provides the standardized access method for data from any network (regional, national, EU, etc.), but does not a priori assure semantic interoperability and compatibility of the data models.

In addition to simply providing access to the underlying data, the SensorSA Cascading SOS service can be used to re-annotate the data on-the-fly before sending it to the requesting client applications.

A third SANY pilot implementation focuses on the feasibility of building 'INSPIRE-ready' service networks based on SensorSA components deployed in Austria. In this case data is offered from all Austrian provinces and the background measurements from the measurement network of the Austrian Environmental Agency:



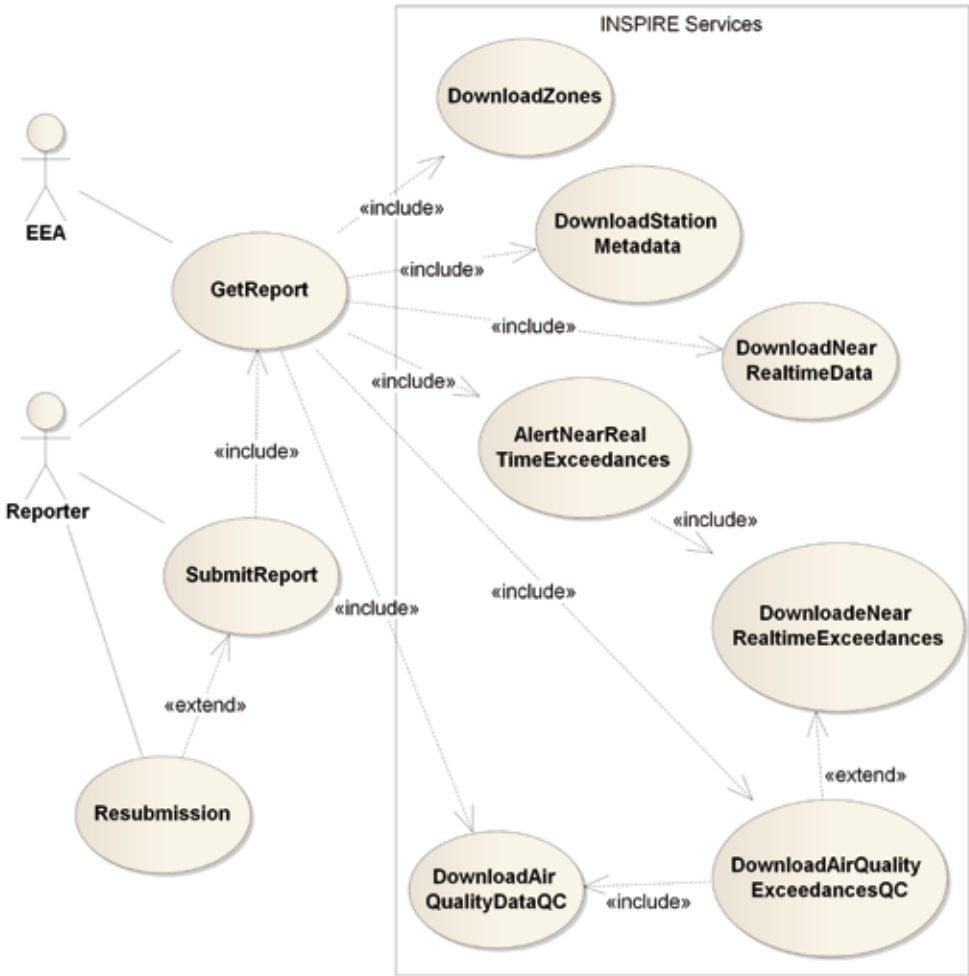
In order to demonstrate the strengths of a decentralized system, the data from two provinces as well as additional background data are provided through separate SOS service instances.

All retrieved data is annotated on-the-fly according to INSPIRE rules for meta-information, and the relevant meta-information is pushed to an INSPIRE-compliant catalogue service.

8.1.4 Report Generation

In addition to providing the INSPIRE-ready metadata model, the pilot also implemented functionality to automate the report generation in order to accommodate periodic national and European reporting obligations.

The following diagram provides an overview of the use cases required for the generation and submission of reports to the European Environment Agency (EEA):



- Within the **DownloadZones** Use Case, the user may download information on the zones and agglomerations defined for assessment and management purposes. This download service will provide all information required for reporting purposes.
- Within the **DownloadStationMetadata** Use Case, the user may download information on the individual air quality monitoring stations. This download service will provide all information required for reporting purposes.

- Within the **DownloadAirQualityDataRaw** Use Case, the user may download raw air quality data as measured by the air quality monitoring stations. This Use Case is closely related to the Use Case DownloadAirQualityDataQC, with the difference that it is not restricted to validated data. This download service is an extension of the Use Case DownloadNearRealtimeData, delivering not only current results, but also aggregating results to the specified time period.
- Within the **DownloadNearRealtimeData** Use Case, the user may download information on the individual air quality monitoring stations.
- Within the **DownloadNearRealtimeExceedances** Use Case, the user may download preliminarily quality controlled exceedance data from air quality monitoring stations, as well as zones and agglomerations, currently showing limit exceedance. This Use Case includes the Use Case DownloadAirQualityDataRaw. This download service will provide all information required for reporting purposes.
- In the **AlertNearRealtimeExceedances** Use Case, alerts are sent in the case of exceedance of thresholds of ozone concentration to a preconfigured application. This Use Case includes the Use Case DownloadNearRealtimeExceedances for the provision of the alert data.
- Within the **DownloadAirQualityDataQC** Use Case, the user may download fully quality controlled data from all monitoring stations. This download service will provide all information required for reporting purposes, with the temporal interval being specified as the required reporting period.
- Within the **DownloadAirQualityExceedancesQC** Use Case, the user may download fully quality controlled exceedance data from all air quality monitoring stations, as well as zones and agglomerations, showing limit exceedance within the given temporal interval. This download service is an extension of the Use Case DownloadAirQualityDataQC, delivering not only current results, but aggregating results for the specified time period. This download service will provide all information required for reporting purposes, with the temporal interval being specified as the required reporting period.

The retrieval and submission process is the same for all reports, but a parameterized data download service has to be provided to support each individual report.

The SANY Cascading SOS concept in conjunction with the Map and Diagram Service provides a suitable means for automatic aggregation and generation of the data required for reporting.

This offers a number of advantages over manual report generation:

- The relevant reporting indicators can be easily reproduced at any time with a minimal effort. This eliminates the main source of errors in report generation.
- The Map and Diagram Service provides a convenient way for automatic generation of maps and diagrams based on the data generated by the Cascading SOS.
- The Cascading SOS and the Map and Diagram Service can be easily used as a back-end for fully automated report generator.

The report data generation is performed by Formula 3 time series processor embedded in the SensorSA Cascading SOS.

Data Plausibility

8.1.5

Data quality assurance can be a tedious, expensive and sometimes also error-prone process that requires continuous supervision of the highly qualified domain experts. Rather than attempting to completely replace the work of domain experts by automatic quality control procedures, SANY looked into options to support domain experts in their work by automatically identifying suspicious measurements.

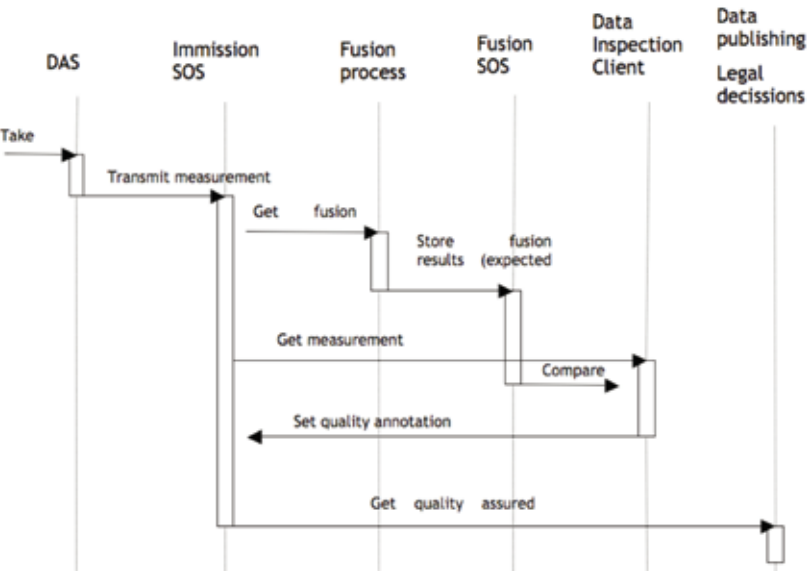
In order to achieve this goal, a state-space fusion service has been developed and deployed in the region of Linz. This service continuously monitors all available immission observations and publishes the nowcasts and 24 hours forecasts at 17 measurement locations using the data model similar to the one used by original immission SOS.

In addition to the nowcasts and forecasts, the state-space fusion also provides the confidence intervals for all estimated values. This allows easy identification of the ‘suspicious’ measurement: a measurement is declared ‘suspicious’ when the difference between data nowcast and actual measurement is larger than the confidence interval advertised by the fusion service.

The identification of the suspicious measurement can be signalled to the domain expert either:

- actively, by rising an alert and sending a notification, e.g. by e-mail or SMS, or
- passively, by providing a visual aid to the expert at the moment he or she is ready to perform the routine data control.

SANY concentrated on the second approach and developed the 'Advanced Data Inspection Tool'. The schematic operation is illustrated in the following figure:



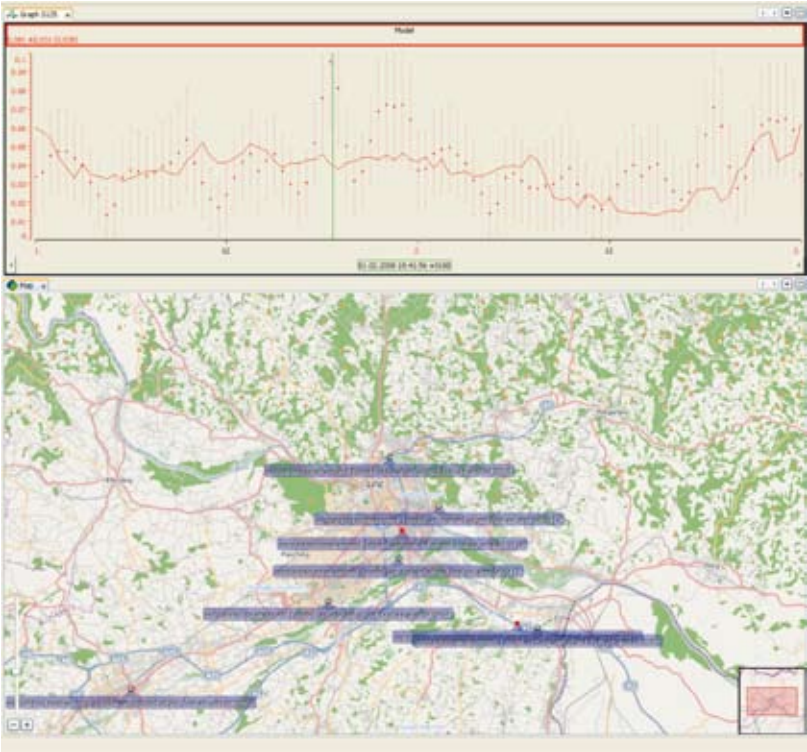
All steps except for the ‘Compare’ and ‘Set quality annotation’ are fully automated. From the users’ point of view, the SANY Data Inspection Tool provides exactly the same type of functionality experts are used to, with only one notable exception: the measurements and their respective errors can be graphically visualized, thus helping the users to easily spot suspicious data. The main visualisation modes are:

- tabular with colour coding of suspicious values;
- x-t diagrams with confidence intervals for the nowcasts
- Geographic map with colour coding of suspicious values

The tabular visualisation is illustrated in this example:

TimeStamp	comparison meas...	comparison meas...	comparison meas...	comparison meas...	comparison meas...	comparison meas...	comparison meas...	comparison meas...
01.02.2008 01:00:00 +0100	0.039 +0.014 (0.020)	0.036 +0.018 (0.020)	0.044 +0.027 (0.047)	0.038 +0.021 (0.040)	0.026 +0.013 (0.028)	0.000 +0.022 (0.040)	0.010 +0.028 (0.040)	0.016 +0.011 (0.026)
01.02.2008 01:30:00 +0100	0.041 +0.017 (0.024)	0.037 +0.019 (0.022)	0.061 +0.029 (0.048)	0.038 +0.022 (0.042)	0.026 +0.013 (0.028)	0.001 +0.022 (0.042)	0.010 +0.028 (0.040)	0.012 +0.012 (0.028)
01.02.2008 02:00:00 +0100	0.040 +0.017 (0.024)	0.036 +0.019 (0.021)	0.062 +0.029 (0.049)	0.039 +0.021 (0.043)	0.026 +0.013 (0.029)	0.002 +0.026 (0.043)	0.011 +0.027 (0.043)	0.012 +0.013 (0.027)
01.02.2008 02:30:00 +0100	0.039 +0.017 (0.024)	0.034 +0.019 (0.021)	0.060 +0.028 (0.041)	0.035 +0.019 (0.038)	0.040 +0.023 (0.048)	0.027 +0.013 (0.029)	0.001 +0.026 (0.048)	0.029 +0.014 (0.029)
01.02.2008 02:00:00 +0100	0.040 +0.017 (0.024)	0.037 +0.019 (0.020)	0.064 +0.029 (0.048)	0.043 +0.023 (0.038)	0.043 +0.023 (0.038)	0.027 +0.023 (0.034)	0.042 +0.022 (0.040)	0.026 +0.014 (0.019)
01.02.2008 02:30:00 +0100	0.036 +0.014 (0.024)	0.024 +0.014 (0.021)	0.043 +0.024 (0.037)	0.040 +0.021 (0.038)	0.033 +0.020 (0.042)	0.024 +0.012 (0.028)	0.023 +0.022 (0.029)	0.032 +0.021 (0.037)
01.02.2008 03:00:00 +0100	0.033 +0.014 (0.020)	0.019 +0.015 (0.021)	0.038 +0.024 (0.032)	0.046 +0.013 (0.039)	0.019 +0.019 (0.044)	0.031 +0.012 (0.027)	0.004 +0.021 (0.029)	0.029 +0.019 (0.019)
01.02.2008 03:30:00 +0100	0.029 +0.015 (0.022)	0.009 +0.014 (0.030)	0.024 +0.023 (0.029)	0.041 +0.019 (0.034)	0.012 +0.013 (0.047)	0.038 +0.011 (0.029)	0.017 +0.021 (0.030)	0.010 +0.019 (0.032)
01.02.2008 04:00:00 +0100	0.036 +0.014 (0.022)	0.006 +0.014 (0.017)	0.022 +0.022 (0.024)	0.032 +0.017 (0.034)	0.020 +0.014 (0.030)	0.038 +0.011 (0.024)	0.042 +0.024 (0.038)	0.031 +0.019 (0.032)
01.02.2008 04:30:00 +0100	0.040 +0.014 (0.020)	0.017 +0.015 (0.019)	0.000 +0.022 (0.021)	0.032 +0.014 (0.030)	0.026 +0.017 (0.037)	0.019 +0.011 (0.022)	0.043 +0.026 (0.032)	0.023 +0.014 (0.029)
01.02.2008 05:00:00 +0100	0.037 +0.017 (0.024)	0.018 +0.018 (0.021)	0.039 +0.024 (0.032)	0.029 +0.019 (0.040)	0.018 +0.010 (0.021)	0.039 +0.029 (0.034)	0.019 +0.019 (0.019)	0.021 +0.011 (0.012)
01.02.2008 05:30:00 +0100	0.037 +0.017 (0.024)	0.018 +0.019 (0.021)	0.087 +0.027 (0.021)	0.041 +0.022 (0.036)	0.021 +0.011 (0.021)	0.038 +0.031 (0.041)	0.029 +0.028 (0.042)	0.027 +0.014 (0.011)
01.02.2008 06:00:00 +0100	0.038 +0.017 (0.030)	0.015 +0.014 (0.025)	0.054 +0.017 (0.038)	0.040 +0.023 (0.030)	0.038 +0.014 (0.029)	0.001 +0.028 (0.040)	0.042 +0.028 (0.039)	0.034 +0.019 (0.013)
01.02.2008 06:30:00 +0100	0.042 +0.014 (0.029)	0.015 +0.017 (0.020)	0.054 +0.026 (0.029)	0.034 +0.019 (0.030)	0.035 +0.021 (0.032)	0.028 +0.015 (0.026)	0.026 +0.024 (0.044)	0.044 +0.019 (0.019)
01.02.2008 07:00:00 +0100	0.038 +0.014 (0.030)	0.017 +0.017 (0.020)	0.087 +0.025 (0.026)	0.039 +0.023 (0.032)	0.027 +0.017 (0.032)	0.024 +0.014 (0.032)	0.024 +0.022 (0.049)	0.036 +0.023 (0.040)
01.02.2008 07:30:00 +0100	0.038 +0.020 (0.034)	0.011 +0.014 (0.020)	0.083 +0.027 (0.031)	0.044 +0.020 (0.034)	0.028 +0.017 (0.030)	0.019 +0.012 (0.027)	0.023 +0.022 (0.031)	0.029 +0.014 (0.014)
01.02.2008 08:00:00 +0100	0.039 +0.020 (0.030)	0.017 +0.015 (0.020)	0.046 +0.020 (0.030)	0.036 +0.021 (0.030)	0.036 +0.021 (0.030)	0.030 +0.023 (0.030)	0.033 +0.014 (0.013)	0.041 +0.014 (0.011)
01.02.2008 09:00:00 +0100	0.038 +0.017 (0.041)	0.027 +0.018 (0.032)	0.051 +0.025 (0.041)	0.045 +0.019 (0.031)	0.046 +0.021 (0.036)	0.038 +0.017 (0.040)	0.044 +0.017 (0.040)	0.044 +0.017 (0.031)

A visualisation using x-t diagrams and mapping is shown here:



Identifying Pollution Impact

8.1.6

Identifying the impact of known pollution sources on actually measured immission provides an indication for the relative importance of pollution sources at selected positions, which are either not known or not taken into account. In other words: whilst the major immission sources tend to be known, additional immissions from background sources will lead to higher measurement levels.

SANY implemented a dispersion modelling service, which takes the real time emission data from all major industrial sites in the city of Linz and meteorological data as input. It calculates the dispersion of the emissions, and produces the prediction of the contamination load at the positions of the immission measurement stations. Thus the estimated immission from known sources are correlated to the actual immission measurements.

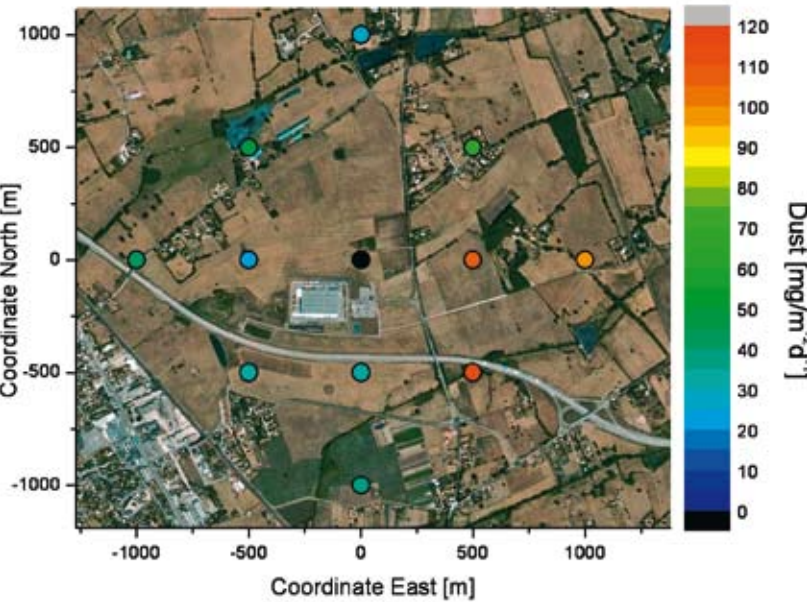
The predictions are published using the SOS service, and the output data model is similar to the one used by immission SOS. This allows easy comparison of the predicted and measured values of immission.

The ratio between the measured and calculated level of pollutants at each measurement point is an indicator for the importance of the ‘background’ sources of pollutions, which could be traffic, households, and emissions that do not originate from the Linz area.


A simplified version of this use case can also be used to evaluate the relative importance of the emission from a single industrial plant to the air quality in its vicinity. In this context, the modelling supplements the immission measurement in two ways:

- First the modelling exposes the influence of new, unknown sources
- Second, the predicted ‘immission from known sources’ is calculated for the whole area, while the measurements are only performed at a small number of points.

This has been tested on the example of an incinerator plant in Moulins, France. The figure below shows the positions of the measurement points around the plant:



The next illustration compares the predicted immission caused by the industrial plant in the middle of the figure and the measured immission.



The comparison clearly shows that the influence of the incinerator plant is negligible at no more than 5 mg per square meter and day, compared to the background immission, which rise up to 120 mg per square meter and day.

DECISION SUPPORT TOOLS FOR MARINE RISK MANAGEMENT

8.2

Forecasting Bathing Water Quality

8.2.1

EU efforts for ensuring clean bathing waters commenced in the 1970s. The 1976 Bathing Water Directive aims at protecting public health and the environment by keeping our coastal and inland bathing waters free from pollution.

What is a bathing water ?

Bathing waters can be coastal waters or inland waters (rivers, lakes). To be covered by the Directive including its mandatory quality standards as well as its monitoring and information obligations, bathing must either be explicitly authorised, or not



prohibited and traditionally practiced by a large number of people. Swimming pools and waters for therapeutic purposes are not covered.¹¹

Bathing beaches often suffer from periodic incidents of reduced water quality due to microbial contamination of bathing waters. These incidents are usually caused by run-off or overloading of urban waste water treatment works after heavy rainfall. Where diffuse pollution sources are the cause, several run-off locations may be suspected, but the actual point of contaminated run-off is likely to vary from incident to incident.

Without a forecasting tool, the decision to close a beach must be based on educated guesswork. In some situations it may be possible to use an in-situ assay technique to measure microbial contamination levels, but results of this technique take 24 hours to obtain. This method is therefore only useful to confirm a contamination incident retrospectively.

Authorities within the EU must be able to forecast contamination incidents in order to meet water quality criteria defined in the EU Bathing Water Directive. To comply with the directive, bathing water samples are taken on pre-specified dates during the bathing season. The table below shows the statutory thresholds for water quality as provided in the Bathing Water Directive:

INLAND WATERS				
Parameter	Excellent	Good	Sufficient	Poor
Intestinal enterococci (cfu/100ml)	200*	400*	330**	< 330
Escherichia coli (cfu/100ml)	500*	1000*	900**	< 900

* Based upon a 95-percentile evaluation

** Based upon a 90-percentile evaluation

¹¹ http://ec.europa.eu/environment/water/water-bathing/index_en.html

COASTAL WATERS AND TRANSITIONAL WATERS				
Parameter	Excellent	Good	Sufficient	Poor
Intestinal enterococci (cfu/100ml)	100*	200*	185**	< 185
Escherichia coli (cfu/100ml)	250*	500*	500**	< 500

* Based upon a 95-percentile evaluation

** Based upon a 90-percentile evaluation

The European Commission frequently publishes information on bathing water quality based on reported water quality measurements and compliance with quality ratings. In order to minimise the health risk on days when threshold exceedance occurs or is expected, the authority must decide by 09:00hrs on every sampling day whether to close the beach in order to ‘discount’ the sample from its annual compliance assessment, which is reflected in the published reports. However, the decision to close a beach can only be based on a forecast of the risk of adverse water quality on the day in question.

The purpose of this SANY Pilot implementation is to provide a forecast tool of water contamination risk. It calculates the probability of contamination exceeding a specified threshold level, at a specified time in the future, for bathing water within the specified region of interest. The bathing water risk application was piloted in the region of the Gulf of Gdansk in Poland. Historical microbial data from water samples taken along the beach at Sopot from January 2001 to December 2007 was combined with meteorological data and the correlation between the two determined using statistical methods for the prediction of the risk of exceedance.

These methods rely on the provision of historical meteorological and microbial data from the chosen bathing water, along with any other available contextual data. This data is then analysed in order to determine the correlation between individual factors and microbial contamination over the historical time period provided. These correlation factors are then used to predict the Risk of Exceedance given real-time values of meteorological and environmental variables.

Three approaches are currently available:

1. Multiple Linear Regression
2. Probabilistic
3. Artificial Neural Network

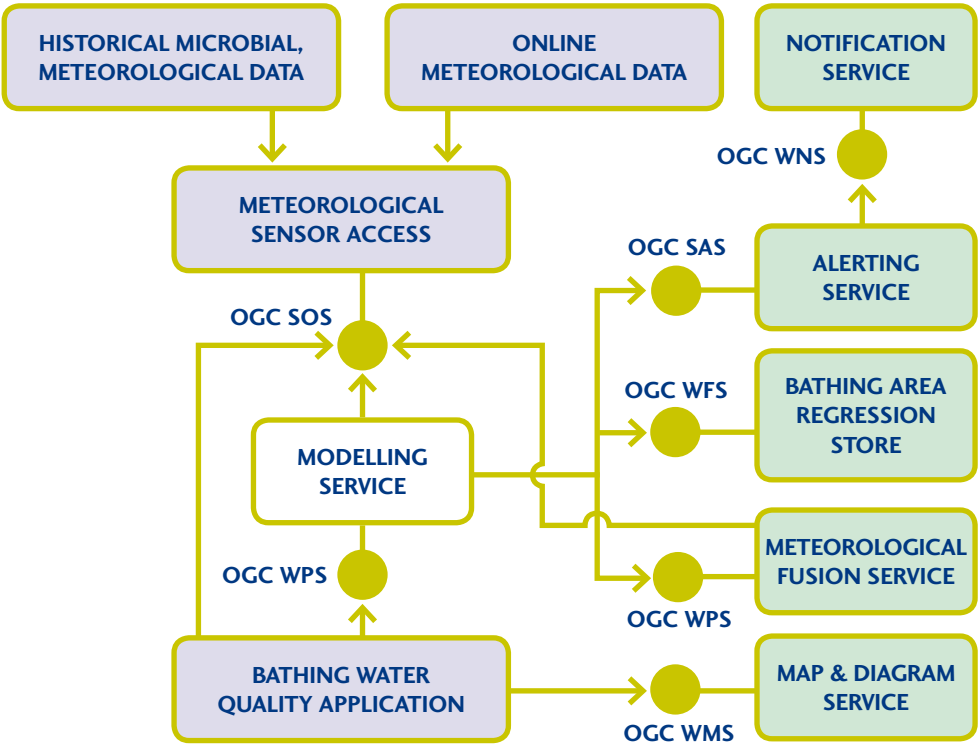
The risk prediction tool developed for the various Polish sites enables the authorities to manage the public hazard of short-term microbial pollution at popular tourist bathing waters within the legislative framework.

The Sensor Service Architecture and decision support tools developed by SANY support the implementation of improved solutions for such situations. For effective bathing water management, it is possible to exploit deployment of additional sensors in selected locations, measuring near-real-time parameters, enabling the forecast of water quality at beaches.

Because of the difficulties associated with the direct measurement microbial contamination levels, risk forecasting requires the monitoring and/or modelling of proxy parameters in order to estimate the contamination risk. Historic measurements of contamination can be used to ‘tune’ these models.

Equally, future measurements of contamination can be used to further validate such models. The application requires access to weather sensors, e.g. rainfall gauges, and, if possible, published weather forecasts. Data fusion is used to populate meteorological data fields across the whole region of interest where the contamination model is being run.

The workflow of the application is shown below:



The supporting architecture for the bathing water risk application has been developed in SANY based on Web services, including a Sensor Observation Service (available as open source implementation) and a Catalogue/Discovery service developed in previous projects.

All external data sources, i.e. sensors, data providers, databases etc. are accessed through the SOS. Data processing, modelling and visualisation are accessed via a range of services, which are registered within the catalogue. All of these services are accessed through the main web based bathing water risk application.

To issue the user with a warning in case of a high risk prediction, a Sensor Alert Service (SAS) and Web Notification Service (WNS) are deployed.

Access to the results and input data for fusion services and the contamination risk model is managed through a Web Feature Service (WFS).

Assessing Forecasts Quality

8.2.2

The quality of risk forecasts for bathing waters is largely reliant on the availability of real-time meteorological data. Since data is not available for all possible locations, spatial interpolation, such as Kriging, may be used to calculate reasonable estimates. As with all estimates, this method introduces some form of error to the data, which needs to be assessed.

The quality of the risk forecast produced from the various statistical approaches can be determined from historical data. An assessment of the accuracy of the prediction has been performed as part of the validation of the statistical methods employed.

The following two tables show the results of validation work carried out for two beaches (Beach A with 16 total pollution events and Beach B with 19 total pollution events). At Beach A, the outlined approaches predict water conditions accurately around 80% of the time. At Beach B, this value drops to 70%.

Validation is based on the parameters learned from the training data. Prediction performance varies from beach to beach depending on data quality, environmental variables available, location, and other unknown factors.

True predictions are cases when both the actual (observed) and predicted bacterial levels fall below the safe threshold OR both the actual (observed) and predicted bacterial levels are above the safe threshold.

BEACH A	MULTIVARIATE LINEAR REGRESSION	PROBABILITY PREDICTION	PROBABILITY MODEL ENSEMBLE (363 MODELS)
NTrue Prediction	72(81.8%)	69(78.4%)	73(83.0%)
NTrue Alarm	10(62.5%)	13(81.3%)	15(93.8%)
NFalse Prediction	16(18.2%)	19(21.5%)	15(17.0%)

- Notes:
- Length of training data is 200 data points from 01/05/1990 to 26/06/1999
 - Length of validation data is 88 data points from 01/06/1999 to 13/09/2004
 - Ensemble approach considers a model set with different configurations of input variables.

BEACH A	MULTIVARIATE LINEAR REGRESSION	PROBABILITY PREDICTION	PROBABILITY MODEL ENSEMBLE (93 MODELS)
NTrue Prediction	50(72.5%)	54(78.3%)	50(72.5%)
NTrue Alarm	10(52.6%)	9(47.4%)	6(31.6%)
NFalse Prediction	19(27.5%)	15(21.7%)	19(27.5%)

- Notes:
- Length of training data is 200 data points from15/05/1990 to 15/06/1999
 - Length of validation data is 69 data points from 19/06/1999 to 17/09/2004
 - Three models accurately predicted the water conditions more than 70% of time at Beach B.

Access to historical microbial data over a period of years is a pre-requisite for the development of the statistical modelling tools, which form the basis of the Bathing Water Risk Management Application. This microbial contamination data is typically collected and stored by the environmental authority of a country or region.

Similarly, historical meteorological and other environmental data for the same period is required and can usually be obtained from the main meteorological provider of a country. There is also a wide range of online weather data sources available, for example, the National Oceanic and Atmospheric Administration (NOAA) in the US, providing access to a worldwide repository of historical weather observations.

Live meteorological, environmental and forecast data, which is suitable input for the risk prediction tool, can be obtained from websites or directly from sensors deployed at the location of the bathing water.

All data that is to be used by the risk prediction tool must be provided through a Sensor Observation Service.

Shellfish Water Quality

In terms of water quality monitoring, shellfish farming has a lot in common with bathing waters: designated shellfish waters also experience intermittent short term microbial pollution, usually associated with high rainfall and flooding events. Shellfish may accumulate microbial contamination from the surrounding environment during these times, thus potentially posing a risk to public health unless the shellfish are treated accordingly. This may include re-locating the shellfish to an uncontaminated environment until any potential contamination has virtually been flushed out.

Shellfish waters are graded on the concentration of E.coli within for example, mussel flesh, as defined by the European Shellfish Waters Directive (2006/113/EC).



8.2.3

The table below shows the statutory thresholds for water quality as outlined in the Shellfish Waters Directive.

Category A	Less than 230 E.coli/100g shellfish flesh	May go for human consumption if End Product Standard* met
Category B	Less than 4,600 E.coli/100g shellfish flesh in 90% of samples	Must be depurated, heat-treated or relayed to meet Category A requirement
Category C	Less than 46,000 E.coli/100g shellfish flesh	Must be relayed for long periods (at least two months) whether or not combined with purification, or after intensive purification to meet Category A or B
	More than 46,000 E.coli/100g shellfish flesh	Unsuitable for production

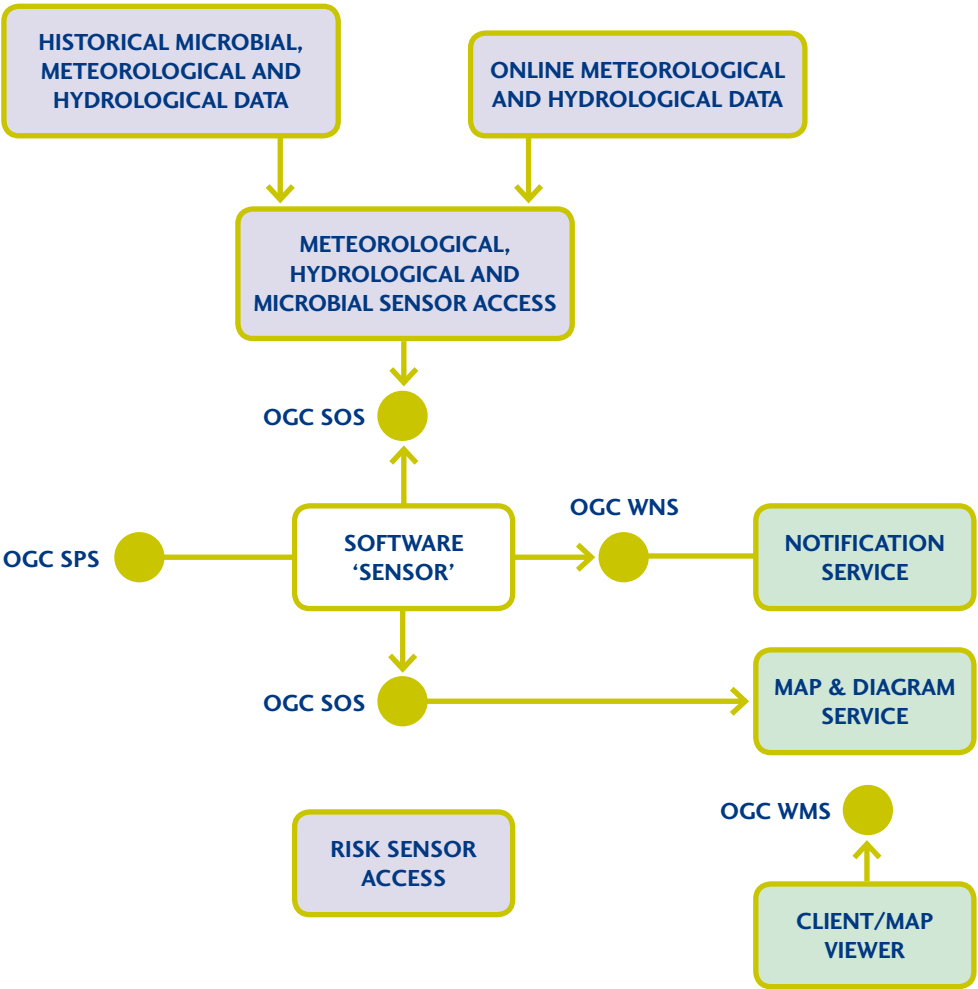
* A requirement to be met before a product can be marketed

Forecasting the onset of short-term microbial pollution events can assist shellfish farmers in managing farming activities, for example, by delaying harvesting when a high risk is forecast, or removing stock stored in racks ahead of an event. This in turn reduces the risk to the consumer.

SANY has developed an application that forecasts the risk of exceedance of the limits for microbial contamination for specified shellfish waters and provides the result directly to decision makers in the aquaculture industry and local authorities. The application is based on OGC compliant services, which include a Hydrology Sensor Observation Service, specialised in retrieving live hydrology parameters from ad-hoc, mobile sensing platforms such as buoys.

As part of the pilot two marine data monitoring buoys were deployed in the Fowey and Fal estuaries in Cornwall, UK, which collect real time hydrographic data from the Fowey and Fal estuaries over a period of several months. This hydrographic data was then combined with meteorological data from the river basins and microbial data from the shellfish waters. The merged data set was used to produce a statistical model for the prediction of short term microbial pollution events in the rivers.

The following figure shows the workflow and the functional building blocks of the Shellfish Water specific application:



The shellfish water application has used the same approach for data modelling, data access and accuracy assessment which has been taken for the Bathing Waters application. This again highlights the flexibility and versatility of deploying service components that support open standards based interfaces to build interoperable solutions.

8.3 GEO-HAZARDS IN URBAN AREAS

Construction work in densely populated urban areas, such as on metros, tunnels, etc. have an impact on the ground as well as on surrounding buildings. Tunnel excavation can induce excavation deformation, surface settlements, and stresses and movements that may lead to severe building damages, landslides, or even the collapse of a tunnel. Those geo-hazards pose the risk of strong economical, environmental and societal impact.

Geo-hazard accidents due to tunnel excavation are very costly: apart from the damages to existing infrastructure, additional work on the tunnel construction can incur severe fines for the delay in completing the construction works. Since those costs can reach millions of Euros, any accident can be very costly!

Since 1992, 54% of tunnels that had no real-time monitoring during construction were damaged, versus 4% for tunnels that were monitored. Therefore a good real-time geotechnical monitoring is essential for the prevention of such accidents.

8.3.1 Aspects of Geotechnical Monitoring

The critical point for the SANY geotechnical pilot application is to both allow the end-user to access all available data and get summed-up or synthetic data to support fast and well-informed decisions. The application must also offer access to enhanced fusion and processing services that will deliver added-value information.

A variety of monitoring aspects needs to be accommodated:

8.3.1.1 Real-time Access to Sensor Data

Since real-time monitoring is quite complex and expensive, the monitoring area during construction works is often limited to the area of active excavation work. This area would move along the tunnel route, following the excavation advancement. The measurements are done in real-time and continuously to immediately identify any changes and trigger response.

Accordingly a high number of sensors needs to be installed and the resulting data volume that needs to be managed can be significant. The system must provide the user with efficient data access, as well as a comprehensive data visualisation; this includes focusing on the most important sensors and areas, and providing clear and easily understandable information.

The system should also allow the user to access the data from any location at the construction site and support seamless access to multiple data sources. This provides the user with a global view of the construction site and helps to prevent fragmentation of information.

In the SANY pilot implementation, the use of a catalogue service allows a user to know which data is available, including detailed information on the monitoring periods, the available sensor systems, etc. and thus allows the user to connect to the appropriate service. The implementation of the Sensor Observation Service (SOS) allows a centralised access to sensor data from different sites, and from different sensor types. The data modelisation used is based on offering concepts and allows the user to define groups of sensors that represent a phenomenon over an area of the site.

The use of two different ResultModels for the SOS allows the user to retrieve:

- times series observations: they show the behaviour of one or several sensor over a defined period, allowing the user to check the trend of sensor(s)
- coverage observations: they show the movements given by all sensor in an area, illustrating the behaviour of soil/structure movements in this zone.

The use of an access control services on top of the SOS ensures that access to sensor data is only granted to authorised users.

Alarms and Alert for an Early Warning

8.3.1.2

The raw data volume generated each day by the monitoring system is too high for an individual person to comprehend and monitor the situation on site. As a decision-maker in charge of a construction site you have little time to analyse each measurement, so a decision support system is needed. This system should provide indicators to raise your attention when needed, but only then. Generally the monitoring system integrates threshold values, which are initially defined and will serve as reference for the definition and issue of alerts. Whenever a threshold is exceeded, an alert is issued and the user will be notified.

The SANY pilot implementation uses a Sensor Alert Service (SAS) that enables the user to define alert threshold and alert conditions over specified sensor measurements. Thus the user can customise alerts according to his specific requirements which will be sent through the Web Notification Service (WNS). The WNS allows the user to define one or several notification means per alert event so that he can choose between different notification means, such as XMPP, Email, SMS, fax, etc. An 'alert panel' lists the history of alerts issued to the user where alert event histories can be checked and deleted.

Estimation of Future Measurements

8.3.1.3

The use of temporal fusion algorithms can be very useful to forecast sensor measurements, and plan remediation works in advance. This kind of fusion is useful to raise the attention on a particular section of the site, and allow a closer

monitoring of the section to prevent potential accidents. The forecasted value is given with a level of confidence or uncertainty.

Temporal fusion may also be used to predict the measurements of a failed sensor, depending on the values of the neighbouring sensors.

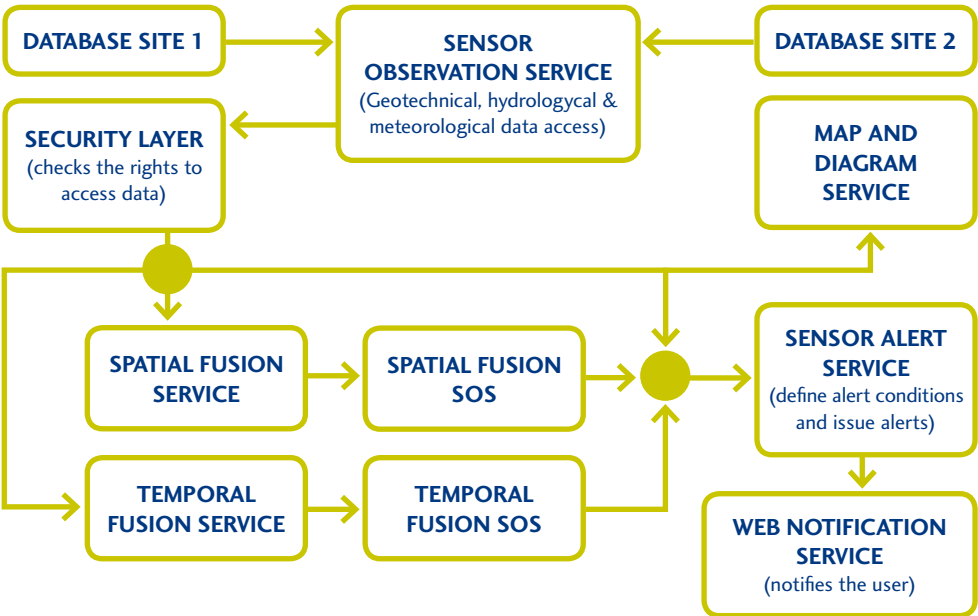
The Web Processing Service for temporal fusion developed in SANY permits the user to get sensor predictions for the next 4h, 6h, 12h, 24h and 48h. The temporal fusion will deliver forecasted values, as well as a level of confidence. According alerts may be issued when the predicted value exceeds a defined threshold.

8.3.1.4 Complement Information at Local Scale

Besides large scale information on the behaviour of individual sensor devices, small scale overview information is also important to understand the overall behaviour of ground and structures in a particular zone. Whilst actual measurements might only be available for a limited area, spatial fusion predictions support the creation of an overview. The results of the spatial fusion algorithms in this case complement the actual available data.

The Web Processing Service for spatial fusion offers gridded information based on sensor measurements. Coupled with a Map & Diagram Service, the visualisation of gridded data as isolines is possible, and those graphical representations may be customised at wish by the user.

The figure below illustrates the workflow implemented in the tunnel construction monitoring scenario.

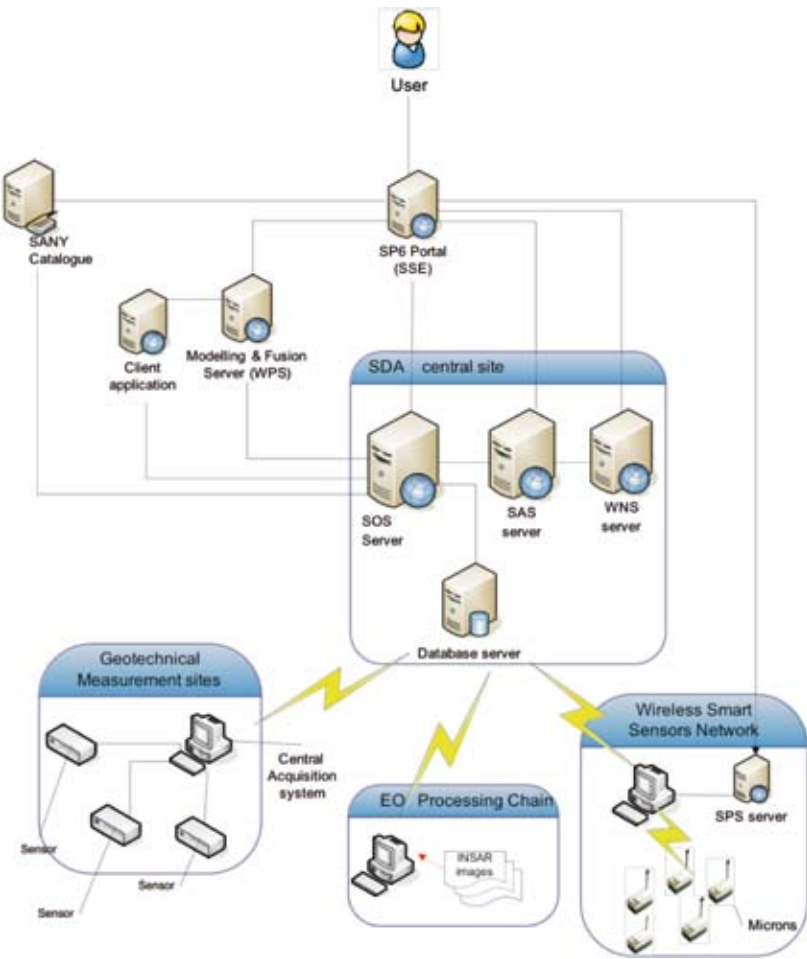


The geo-hazard pilot developed within SANY covers the following applications:

- monitoring of an area around a tunnel construction site;
- management of sensor networks for geo-hazards;
- management of remediation techniques.

The pilot implementation is based on the Sensor Service Architecture and all external data sources of are accessed through the Sensor Observation Service. Data processing, modelling and visualisation are accessed via a range of services, which are registered within the catalogue.

The complete overview of this application is shown below:



8.3.1 Management of Sensor Networks

In the case of ‘classical’ geotechnical sensor systems, the sensors are initially installed, configured and connected to wired data loggers and communication networks, and their installation requires a lot of preparation and time. Likewise maintenance and configuration changes require staff to be on site and the multiplication of such operations is time-consuming and may even present a risk for the operation staff.

In case of hazard or suspicion of hazard, it is important to get information where needed, in a timely fashion way. Geotechnical sensor networks are fixed, wired, and hardly flexible. The addition of new sensor devices requires time and is not done in an easy way. Therefore, especially in case of crisis, the need for sensors that can be fast and easily deployed, and which require a minimum configuration is essential. Devices must be quickly deployed in order to reduce staff’s exposure to any risk, and the communication between the devices must be flexible and dynamic for the data communication to be ensured even when the situation on site change. Moreover, the new data issued by those sensors must be integrated to the existing sensor system, allowing to complement efficiently existing data.

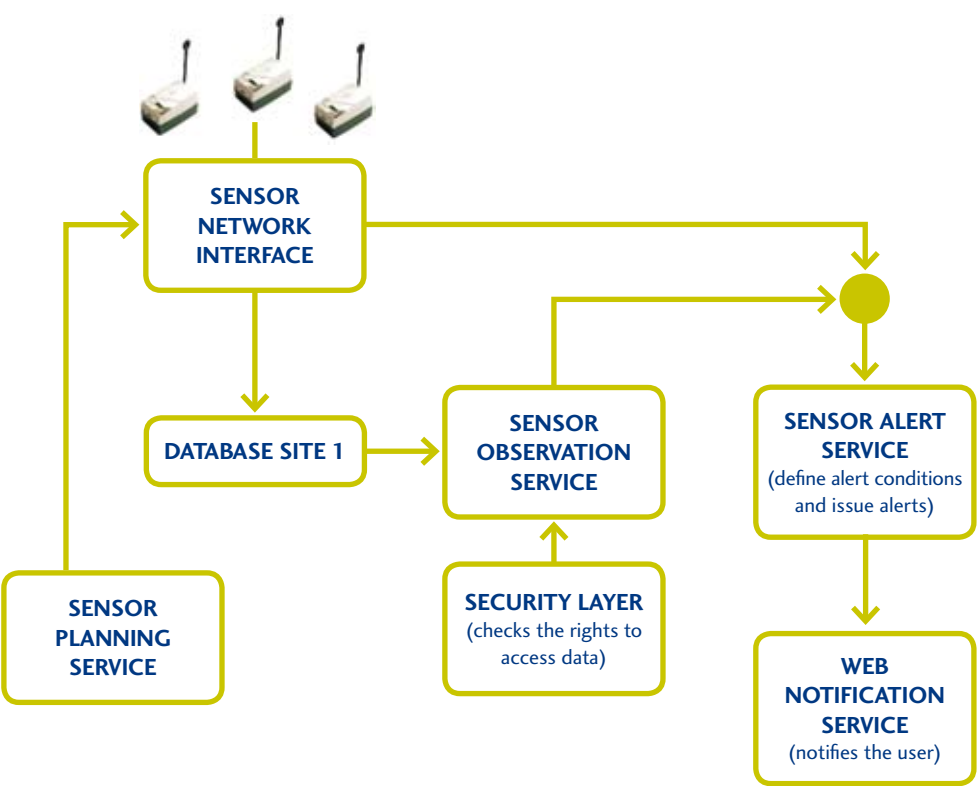
It is therefore advantageous to have a service which facilitates remote access to the sensor configurations, the battery level, and if needed to remotely configure the sensor device.

The maintenance of sensor systems may be eased and gain in efficiency through the provision of information about a sensor failure and sensor battery level. An alarm system based on the battery level of the sensors or the failure would allow the operator to be warned before a node runs out of battery (so he can plan in advance maintenance operation, order a battery, etc.) or in case of sensor failure over a defined period, the operator would be notified and send someone on site to check and replace the sensor if needed.

SANY has developed smart sensor nodes which are organised in a flexible wireless network, are autonomous, and connect to a wide variety of sensors. Therefore the extension of existing sensor networks and addition of new measurement points is made easy and efficient. As those devices are autonomous and work on a self-healing and self-organising network, their configuration is reduced to the minimum, and they can be left on site without the need of intervention, even when the site configuration changes.

With the use of a Sensor Planning Service, coupled with those devices, the reconfiguration of the devices can be done remotely. It is also possible to check the configuration (battery level, radio configuration and power) through the service, reducing the need to go on site for maintenance operations. Additional smart sensor nodes can be added in a ‘plug and play’ manner to the network, and their information will be automatically integrated into the system.

The acquisitions made by the sensors are stored in a Sensor Observation Service, and can be accessed with the SOS client. The workflow is illustrated in the figure below.



Management of Remediation Techniques

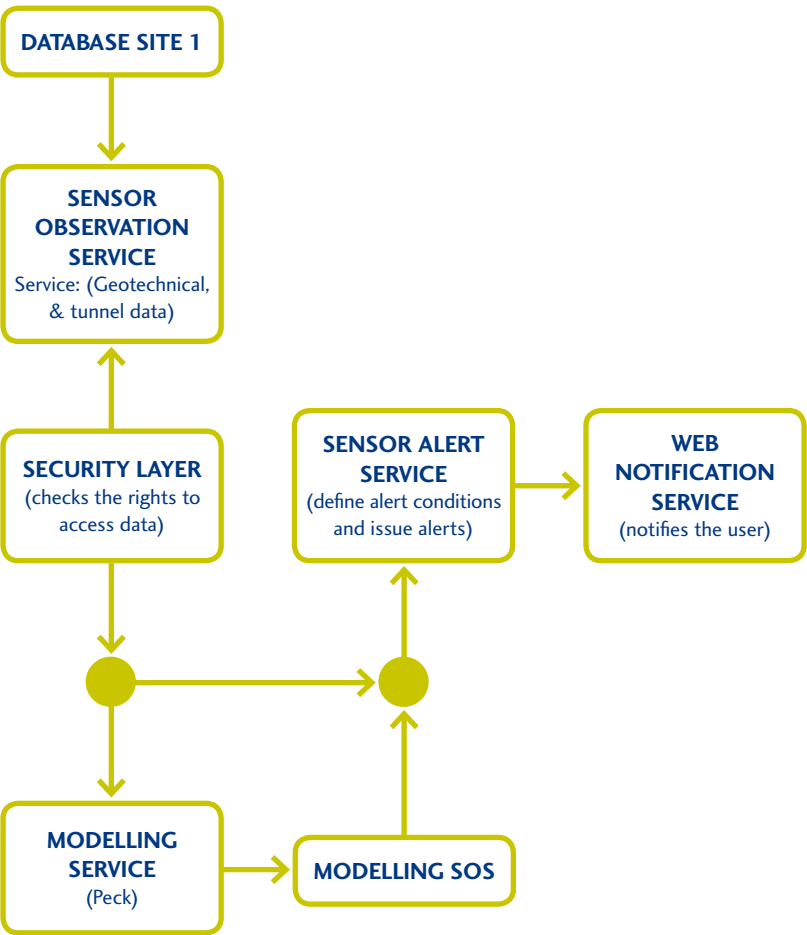
8.3.2

If significant settlements are detected, remediation techniques are required to prevent major consequences. Efforts to control ground losses and settlement can be made by improving granular soils by grouting/replacement techniques or by using specially modified earth pressure balance or slurry pressure balance tunnelling machines; however, these techniques are not always successful. To manage these remediation techniques in the best way, settlement predictions must be available in order to adjust the technique to the specific site parameters.

The development of a settlement prediction model based on in-situ real-time data would allow finer modelling and more reliable information. The SANY geotechnical application has combined a Peck model with sensor data, so that its parameters are refined and adapted to the real site conditions over time. The Web Processing Service that runs the Peck model takes information on

the tunnel and sensor measurements as input: the real sensor measurements feed the model, which returns a prediction of final settlements at the point of interest. The major advantage of this approach is that it is not only based on a pre-defined model, but it is constantly refined with real measurements, thus offering more accurate predictions, which fit well with real observations.

The service based on this model also allows a comparison of predicted final settlements with the current trend and initial theoretical values. If the settlements predicted are higher than expected, an alert is triggered by coupling this service with a Sensor Alert Service. The entire workflow is illustrated below.



CONCLUSIONS

8.4

The SensorSA architecture and OGC SWE specifications provide a viable alternative to the proprietary monitoring systems, with the major added value of enhanced interoperability. The SANY pilot implementations clearly show the value of the service oriented approach:

Web processing services, such as fusion and modelling, can be easily added to an existing network, with no significant changes to existing infrastructure. In fact, several modelling and/or fusion services could be easily run side by side for evaluation purposes with no adverse effects to the normal network operation.

The presentation of measurements and analysis results in a usable form for the decision makers does not require huge efforts in terms of the client software development. Simple presentation can be done using any off-the shelf SOS client. More complex indicators can be calculated on the fly and presented with the help of generic software components which were developed in SANY, such as the Map and Diagram Service, GTV client, the Cascading SOS etc. In fact one of the major advantages of adopting an open standards based approach is the availability of numerous service components, which are available under open source licenses.

In addition to providing a standards compliant wrapping layer for existing monitoring networks, the SANY approach demonstrated how to add new sensors to existing networks or even build complete monitoring networks without proprietary components.

In a nutshell, by making use of the SANY Sensor Service Architecture and adopting open standards based interfaces, the resulting applications will benefit in numerous ways:

- The SANY Catalogue integrates semantic and ontology features and enables the discovery of available resources and services;
- Customisable applications can be based on generic software components and building blocks which can be easily re-configured and used in new application areas;
- The standards based interfaces enable common access to a wide variety of data sources and sensor information;
- The easy and fast deployment and integration of new sensor networks in a 'plug and measure' manner, as well as their combination with other SANY-compatible sensors, enables flexible and responsive monitoring in high risk areas;

- SANY smart sensor nodes are autonomous and can be configured remotely without having to go on site;
- Different sensor data sources can be combined, taking into account their information to evaluate the reliability of the resulting information;
- Spatial fusion services developed by SANY provide richer information;
- Measurements and resulting information can be presented with enhanced visualisation services based on customisable configuration parameters;
- The possibility for the user to define its own alert conditions, and, in combination with the notification service, set-up an early warning system;
- Standard descriptions of sensors allow them to be shared and used by different services and application domains.

Last, not least, the work of SANY with end users in the context of the pilot implementations has provided the consortium partners with a wealth of experiences in using standards based components. It has also provided valuable feedback, which is currently fed back into the standardisation process to improve existing specifications wherever gaps have been identified. The SANY consortium partners have a strong interest to share the knowledge that has been gained in the project and help to promote, as the acronym suggest, the establishment of SENSORS ANYWHERE.

If you require more information and want to learn more about Sensor Networks, get in touch! info@sany-ip.eu

Annex 1: Abbreviations and Acronyms

ADC Architecture and Data Committee (GEOSS)	GFM General Feature Model
CAFE Clean Air for Europe programme	GML Geographiv Markup Language
DCP Distrinuted Computing Platform	GMES Global Monitoring for Environment and Security
DG-INFSO Directorate General for Information Society (EC)	HMA Heterogeneous Missions Accessibility
DoW Description of Work	HTTP Hypertext Transfer Protocol
EC European Commission	ID Identifier
ECMWF European Center for Medium range Weather Forecasting	IEEE Institute of Electrical and Electronics Engineers
EO Earth Observation	IETF Internet Engineering Task Force
ESA European Space Agency	INSPIRE Infrastructure for Spatial Information in Europe; Framework directive for building an infrastructure for spatial information in the Community (www.inspire.jrc.it)
ESDI European Spatial Data Infrastructure	IS International Standard
EU European Union	ISO The International Organization for Standardization
FOI Feature of Interest	IST Information Society Technology
FP6/7 6th/7th Framework Programme (EC)	IT Information Technologies
GEOSS Global Earth Observation System of Systems	JRC DG Joint Research Centre (EC)

LDAP
Lightweight Directory Access Protocol

MIB
Management Information Base

OASIS
Organization for the Advancement
of Structured Information Standards
(www.oasis-open.org)

OCC
Open Geospatial Consortium
(www.opengeospatial.org)

OMG
Object Management Group

ORCHESTRA
Open Architecture and Spatial Data
Infrastructure for Risk Management
(FP6 project)

ORM
OGC Reference Model

OWS
OGC Open Web Services Testbed

O&M
Observations and Measurement

QoS
Quality of Service

PDP
Policy Decision Point

PEP
Policy Enforcement Point

PIP
Policy Information Point

RDF
Resource Description Framework

REST
Representational State Transfer

RM-OA
Reference Model for the ORCHESTRA
Architecture

RM-ODP
Reference Model for Open Distributed
Processing

SAML
Security Assertion Markup Language

SANY
Sensors Anywhere (FP6 project)

SAS
Sensor Alert Service

SAWSDL
Semantic Annotation for WSDL and XML Schema

SDI
Spatial Data Infrastructure

SensorSA
Sensor Service Architecture

SIF
Standards and Interoperability Forum (GEOSS)

SLD
Styled Layer Descriptor

SOA
Service-oriented Architecture

SOAP
Lightweight protocol to exchange
xml-based messages

SOA-RA
(OASIS) Reference Architecture for Service
Oriented Architecture

SOA-RM
(OASIS) Reference Model for Service
Oriented Architecture

SOS
Sensor Observation Service

SPS
Sensor Planning Service

SSE
Service Support Environment

SWE
Sensor Web Enablement

UAA
User Management, Authentication
and Authorisation

UDDI
Universal Description, Discovery
and Integration

URI
Uniform Resource Identifier

URN
Uniform Resource Name

UTC
Universal Coordinated Time

WADL
Web Application Description Language

W3C
World Wide Web Consortium (www.w3.org)

WFS
Web Feature Service

WMS
Web Map Service

WNS
Web Notification Service

WPS
Web Processing Service

WSDL
Web Service Description Language

XML
eXtensible Markup Language

XACML
eXtensible Access Control Markup Language

Annex 2: Glossary

Absolute Time Provides 1) a means to specify an absolute time (UTC) for meta-information, and 2) a general-purpose mechanism for describing points in absolute (UTC) time. (derived from ISO/IEC 18023:2006(E))	Application Use of capabilities, including hardware, software and data, provided by an information system specific to the satisfaction of a set of user requirements in a given application domain. (derived from OGC glossary)
Access control Ability to enforce a policy that identifies permissible actions on a particular resource by a particular subject.	Application Domain Integrated set of problems, terms, information and tasks of a specific thematic domain that an application (e.g. an information system or a set of information systems) has to cope with. An example of an application domain is environmental risk management.
Accounting Process of gathering information about the usage of resources by subjects.	Application Schema Conceptual schema for data required by one or more applications. (ISO 19109:2005)
Ad hoc Sensor Network Sensor network in which communication links and/or nodes are not continually available or change dynamically. An ad hoc sensor network is often, but not necessarily, based on wireless communication between nodes with limited resources (energy supply, processing power). An ad hoc sensor network may include mobile sensors which belong to the network for a limited time or intermittently.	Application Architecture Instantiation of a generic and open architecture by inclusion of those thematic aspects that fulfil the purpose and objectives of a given application. The concepts for such an application stem from a particular application domain.
Alert Message that transports one or more events. Depending on the form of the event, the notification may resemble the event that it transports. <i>NOTE: Used as a synonym for notification.</i>	Architecture (of a system) Set of rules to define the structure of a system and the interrelationships between its parts. (ISO/IEC 10746-2:1996)
	Architecture Service Service that provides a generic, platform-neutral and application-domain independent functionality.

Authentication Concerns the identity of the participants in an exchange. Authentication refers to the means by which one participant can be assured of the identity of other participants. (SOA-RA)	Conceptual Schema The formal description of a conceptual model. (ISO 19109:2005(E); ISO 19101)
Authorisation Concerns the legitimacy of the interaction. Authorisation refers to the means by which an owner of a resource may be assured that the information and actions that are exchanged are either explicitly or implicitly approved. (SOA-RA)	Confidentiality Concerns the protection of privacy of participants in their interactions. Confidentiality refers to the assurance that unauthorized entities are not able to read messages or parts of messages that are transmitted. (SOA-RA)
Catalogue Collection of entries, each of which describes and points to a collection of resources. Catalogues include indexed listings of resource collections, their contents, their coverages, and of meta-information. A catalogue registers the existence, location, and description of resource collections held by an Information Community. Catalogues provide the capability to add, modify and delete entries. A minimum Catalogue will include the name for the resource collection and the locational handle that specifies where these data may be found. Each catalogue is unique to its Information Community. (derived from OGC glossary)	Discovery Act of locating a machine-processable description of a resource that may have been previously unknown and that meets certain functional criteria. It involves matching a set of functional and other criteria with a set of resource descriptions. (derived from W3C glossary)
Component A component can either be a hardware component (device) or software component.	End User Members of agencies (e.g. civil or environmental protection agencies) or private companies that are involved in an application domain (e.g. risk management) and that use the applications built by the system users.
Conceptual Model Model that defines concepts of a universe of discourse whereby the universe of discourse comprises the extract of the real or hypothetical world that includes everything of interest for a particular application (ISO 19109:2005(E); ISO 19101)	Event Anything that happens or is contemplated as happening at an instant or over an interval of time. (derived from ISO 19136)
	Environment 1: (noun) the surroundings or conditions in which a person, animal, or plant lives or operates. 2: (the environment) the natural world, especially as affected by human activity. 3: (computing) Overall structure within which a user, computer, or program operates. (derived from The Oxford Dictionary)

Feature

Abstraction of a real world phenomenon (ISO 19101) perceived in the context of an application. (derived from ISO 19101)

NOTE: The SANY understanding of a ‘real world’ explicitly comprises hypothetical worlds. Features may but need not contain geospatial properties. In this general sense, a feature corresponds to an ‘object’ in analysis and design models.

Framework

An information architecture that comprises, in terms of software design, a reusable software template, or skeleton, from which key enabling and supporting services can be selected, configured and integrated with application code. (derived from OGC glossary)

Generic (Service, Infrastructure...)

Independent on the organisation structure and application domain, etc. For example, a service is generic, if it is independent of the application domain. A service infrastructure is generic, if it is independent of the application domain and if it can adapt to different organisational structures at different sites, without programming (ideally).

Geospatial

Referring to a location relative to the Earth’s surface. ‘Geospatial’ is more precise in many geographic information system contexts than ‘geographic,’ because geospatial information is often used in ways that do not involve a graphic representation, or map, of the information. (OGC glossary)

Identity

Concept that is used to recognise a subject. A subject may have several identities.

Implementation

Software package that conforms to a standard or specification. A specific instance of a more generally defined system. (<http://www.opengeospatial.org/resources/?page=glossary>)

Integrity

Concerns the protection of information that is exchanged – either from unauthorized writing or inadvertent corruption. Integrity refers to the assurance that information that has been exchanged has not been altered. (SOA-RA)

Interface

Named set of operations that characterize the behaviour of an entity. The aggregation of operations in an interface, and the definition of interface, shall be for the purpose of software reusability. The specification of an interface shall include a static portion that includes definition of the operations. The specification of an interface shall include a dynamic portion that includes any restrictions on the order of invoking the operations. (ISO 19119:2005)

Interoperability

Capability to communicate, execute programs, or transfer data among various functional units in a manner that require the user to have little or no knowledge of the unique characteristics of those units (ISO 2382-1). (ISO 19119:2005)

Loose Coupling

Coupling is the dependency between interacting systems. This dependency can be decomposed into real dependency and artificial dependency: Real dependency is the set of features or services that a system consumes from other systems. The real dependency always exists and cannot be reduced. Artificial dependency is the set of factors that a system has to comply with in order to consume the features or services provided by other systems. Typical artificial dependency factors are language dependency, platform dependency, API dependency, etc. Artificial dependency always exists, but it or its cost can be reduced. Loose coupling describes the configuration in which artificial dependency has been reduced to the minimum. (W3C glossary)

Meta-information

Descriptive information about resources in the universe of discourse. Its structure is given by a meta-information model depending on a particular purpose.

NOTE: A resource by itself does not necessarily need meta-information. The need for meta-information arises from additional tasks or a particular purpose (like catalogue organisation), where many different resources (services and data objects) must be handled by common methods and therefore have to have/get common attributes and descriptions (like a location or the classification of a book in a library).

Meta-information Model

Implementation of a conceptual model for meta-information.

Non-repudiation

Concerns the accountability of participants. To foster trust in the performance of a system used to conduct shared activities it is important that the participants are not able to later deny their actions: to repudiate them. Non-repudiation refers to the means by which a participant may not, at a later time, successfully deny having participated in the interaction or having performed the actions as reported by other participants. (SOA-RA)

Notification

Message that transports one or more events. Depending on the form of the event, the notification may resemble the event that it transports.

NOTE: Used as synonym for alert.

Observed Property

Identifier or description of the phenomenon for which the observation result provides an estimate of its value. (derived from OGC 07-022r1)

Observation

Act of observing a property or phenomenon, with the goal of producing an estimate of the value of the property. (OGC 07-022)

Open Architecture

Architecture whose specifications are published and made freely available to interested vendors and users with a view of widespread adoption of the architecture. An open architecture makes use of existing standards where appropriate and possible and otherwise contributes to the evolution of relevant new standards.

Operation

Specification of a transformation or query that an object may be called to execute. An operation has a name and a list of parameters. (ISO 19119:2005)

ORCHESTRA Architecture

Open architecture that comprises the combined generic and platform-neutral specification of the information and service viewpoint as part of the ORCHESTRA Reference Model.

ORCHESTRA Reference Model

The ORCHESTRA Reference Model comprises a specification of all RM-ODP viewpoints for the open architecture for risk management. (<http://www.eu-orchestra.org>)
NOTE: The ORCHESTRA Reference Model is specified in (Usländer (ed.), 2007) and is the result of the European Integrated project ORCHESTRA. The relationship of the SANY Sensor Service Specification to the ORCHESTRA Reference Model is specified in chapter 6 of this book.

Phenomenon

Concept that is a characteristic of one or more feature types, the value for which may be estimated by application of some procedure in an observation. (OGC 07-022)

Plug-and-measure

Refers to the degree of capability to add a new sensor to a sensor network, register it in a sensor service network and access its observations through sensor services in all functional domains of a sensor service network without additional manual intervention.

Policy

Representation of a constraint or condition on the use, deployment, or description of a resource. (derived from SOA-RM)

Purpose (of meta-information)

Describes the goal of the usage of the resources. (OGC 07-097; RM-OA 2007)

(Service) Platform

Set of infrastructural methods, technologies and rules that describe how to specify service interfaces and related information and how to invoke services in a distributed system. Examples for platforms are Web Services according to the W3C specifications including a GML profile for the representation of geographic information.

Reference Model

Framework for understanding significant relationships among the entities of some environment, and for the development of consistent standards or specifications supporting that environment. A reference model is based on a small number of unifying concepts and may be used as a basis for education and explaining standards to a non-specialist. (ISO Archiving Standards; <http://ssdoo.gsfc.nasa.gov/nost/isoas/us04/defn.html>)

Representation

Comprises any useful information about the current state of a resource. (Richardson/ Ruby 2007)

Resource

Anything that's important enough to be referenced as a thing itself. (Richardson/ Ruby 2007)
NOTE: Applied to geospatial service-oriented architectures: Functions (possibly provided through services) or data objects (possibly modelled as features).

Sensor

Entity that provides information about an observed property at its output. A sensor uses a combination of physical, chemical or biological means in order to estimate the underlying observed property. At the end of the measuring chain electronic devices produce signals to be processed.
NOTE: A more detailed discussion about simple and complex forms of a sensor as well as sensor systems, also in the context of environmental models, is given in chapter 6.

Sensor Network

A collection of sensors and processing nodes, in which information on properties observed by the sensors may be transferred and processed.
NOTE: A particular type of a sensor network is an ad hoc sensor network.

Sensor Service Architecture (SensorSA)
Open Architecture comprising a platform-neutral conceptual specification of the architectural components of a service network that includes the access to sensors, sensor networks and sensor-related information.

Sensor System

System whose components are sensors. A sensor system as a whole may itself be referred to as a sensor with an own management and sensor output interface. In addition, the components of a sensor system are individually addressable.

Service

Distinct part of the functionality that is provided by an entity through interfaces. (ISO 19119:2005)

Service Instance
Executing manifestation of a software component that provides an external interface of a service according to an implementation specification for a given platform.

Service Network
Set of networked hardware components and service instances that interact in order to serve the objectives of applications. The basic unit within a service network for the provision of functions are the service instances.

Session
Temporarily valid ticket.

Signal
Any internal representation (i.e. internal to the sensor) of the observed property.

Software Component
Program unit that performs one or more functions and that communicates and interoperates with other components through common interfaces. (derived from OGC glossary)

Spatial Context

Specification of a spatial location of an observed property determined by a combination of a point, a line, an area, a volume and/or a vector field.

NOTE: As an example for the combination of an area and a point, consider a sensor that is capable of recording an image of an area. It may deliver both a spatial context for the area (e.g. the polygon of the area) and/or for several points within that area (e.g. a grid laid upon the area).

Subject

Abstract representation of a user or a software component in an application.

System

Something of interest as a whole or as comprised of parts. Therefore a system may be referred to as an entity. A component of a system may itself be a system, in which case it may be called a subsystem. (ISO/IEC 10746-2:1996)

NOTE: For modelling purposes, the concept of system is understood in its general, system-theoretic sense. The term 'system' can refer to an information processing system but can also be applied more generally.

System User

Provider of services that are used for an application domain as well as IT architects, system developers, integrators and administrators that conceive, develop, deploy and run applications for an application domain.

Temporal Context

Specification of the temporal reference of an observed property based on the absolute time. It can be a single point in time, a time sequence, a time period or a combination of these. In a sampling system for example several time periods and time points are needed to describe the time behaviour. However, a time point is already an abstraction which does not really exist. It means a very small time interval.

Ticket

Information issued by an identity provider to be used as proof of identity when accessing a resource.

Uncertainty

Quantified description of the doubt about the measurement result.
NOTE: The error of a measurement may be small, even though the uncertainty is large.

Universe of discourse

View of the real or hypothetical world that includes everything of interest. (ISO 19101)

Annex 3: Bibliography

Botts, M. 2007. OpenGIS® Sensor Model Language (SensorML) Implementation Specification. OGC® 07-000.
http://portal.opengeospatial.org/files/?artifact_id=21273

Botts, Percivall, Reed & Davidson, 2007. OGC® Sensor Web Enablement: Overview And High Level Architecture. OGC White Paper OGC 07-165.
http://portal.opengeospatial.org/files/?artifact_id=25562

CEHIS consortium 2008. Connectivity between Environment and Health Information Systems. Final Study Report. Version 2.3, SMART 2006/0054, <http://envihealth.jrc.cec.eu.int/CEHIS>

Christakos, G. 2000. Modern spatiotemporal geostatistics. Oxford University Press, New York, 2nd edition (2001), ISBN 0-19-513895-3.

Christakos, Bogaert, & Serre 2002. Temporal GIS: advanced functions for field-based applications. Springer-Verlag, ISBN 3-54041 - 476-2.

Cox, S. (Ed.), 2007. Observations and Measurements – Part 1 – Observation schema. OGC 07-022r1.
<http://www.opengeospatial.org/standards/om>

Cox, S. (Ed.), 2007. Observations and Measurements – Part 2 – Sampling Features. OGC 07-002r3.
<http://www.opengeospatial.org/standards/om>

Duschene & Sonnet (Ed.), 2005. WMS Change Request: – Support for WSDL and SOAP, OGC 04-050r1.
<http://www.opengeospatial.org/standards/wms>

Erl, T. (2008). SOA: Principles of Service Design. Prentice Hall. ISBN 0-13-234482-3.

Fensel et al. 2006. ‘Enabling Semantic Web Services – The Web Service Modeling Ontology’. Springer-Verlag Berlin Heidelberg New York, ISBN-13 978-3-54034-519-1.

Fielding, R.T. 2000. Architectural Styles and the Design of Network-Based Software Architectures. Doctoral dissertation, University of California, Irvine.
<http://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm>

Graham, Hull & Murray (Ed.), 2006. Web Service Base Notification 1.3, OASIS identifier wsn-ws_base_notification-1.3-spec-os.
http://docs.oasis-open.org/wsn/wsn-ws_base_notification-1.3-spec-os.pdf

Handley, Jacobson & Perkins 2006. Session Description Protocol. The Internet Engineering Task Force, IETF RFC 4566.
<http://www.ietf.org/rfc/rfc4566.txt>

Havens, S. (Ed.) 2007. OpenGIS® Transducer Markup Language (TML) Implementation Specification. OGC 06-010r6.
<http://www.opengeospatial.org/standards/tml>

Intamap – Williams, Cornford, Bastin & Pebesma, 2007. Uncertainty Markup Language (UncertML).

<http://www.intamap.org/uncertml/uncertml.php>

ISO GUM 1993. Guide to the Expression of Uncertainty in Measurement. Editors: BIPM, IEC, IFCC, ISO, IUPAC, IUPAP, OIML. International Organisation for Standardization, Geneva, Switzerland. ISBN 92-67-10188-9, First Edition 1993.

Klopper & Kannellopoulos, (Ed.), 2008.

ORCHESTRA – an open service architecture for risk management.

<http://www.eu-orchestra.org/docs/ORCHESTRA-Book.pdf>

Kunz, Usländer, & Watson 2009. A Testbed for Sensor Service Networks and the Fusion SOS: towards plug & measure in sensor networks for environmental monitoring with OGC standards. 18th World IMACS/MODSIM Congress, Cairns, Australia 13-17 July 2009, <http://mssanz.org.au/modsim09>

Lalonde, W. (Ed) 2002. Styled Layer Descriptor Implementation Specification, OGC 02-070. <http://www.opengeospatial.org/standards/sld>

Müller & MacGill (Ed.) 2005. Styled Layer Descriptor Profile of the Web Map Service Implementation Specification, OpenGIS® Discussion Paper 05-078. <http://www.opengeospatial.org/standards/sld>

Müller (Ed.) 2006. Symbology Encoding Descriptor Implementation Specification. OGC 05-077r4.

<http://www.opengeospatial.org/standards/symbol>

Na & Priest (Ed.), 2007. Sensor Observation Service. OGC 06-009r6.

<http://www.opengeospatial.org/standards/sos>

OASIS 2006. Security Assertion Markup Language (SAML) V2.0 Technical Overview Working Draft:

<http://www.oasis-open.org/committees/download.php/20645/sstc-saml-tech-overview-2%200-draft-10.pdf>

Simonis, I. (Ed.), 2007. OpenGIS® Sensor Planning Service Implementation Specification. OGC® 07-014r3.

<http://www.opengeospatial.org/standards/sps>

Simonis, I. (ed.) 2008. OGC® Sensor Web Enablement Architecture, OGC Best Practices Document 06-021r4 Version: 0.4.0, <http://www.opengeospatial.org/standards/bp>, 2008.

Simonis & Wytzisk (Ed.) 2003. Web Notification Service. OGC 03-008r2. http://portal.opengeospatial.org/files/?artifact_id=1367

Percivall, G. (Ed.) 2008. OGC Reference Model. OGC 08-062r4. <http://www.opengeospatial.org/standards/orf>

Richardson & Ruby 2007. RESTful Web Services. O'Reilly Media, Inc. ISBN-10: 0-596-52926-0.

Sabeur Z. A. 2007. Sensors Fusion Services for Environmental Decision-Support. Presented at 'Insitu monitoring and Earth observation in the framework of GMES and GEOSS', EU Commission-EARS conference, Brussels, May 31st 2007.

Usländer, T. (Ed.), 2007. Reference Model for the ORCHESTRA Architecture Version 2 (Rev 2.1), Deliverable D3.2.3 ORCHESTRA FP6 Integrated project 511678, OGC Best Practices Document 07-097.

http://portal.opengeospatial.org/files/?artifact_id=23286

Usländer, T. (ed.) 2009. 'Specification of the Sensor Service Architecture Version 3, SANY Deliverable D2.3.5.

<http://www.sany-ip.eu>.

World Wide Web Consortium (W3C) 2006. Web Services Description Language (WSDL) Version 2.0 Part 1: Core Language, W3C Candidate Recommendation. <http://www.w3.org/TR/2006/CR-wsdl20-20060327>

The SANY Consortium Partners, their acronyms in the project and contact details for SANY:

AIT Austrian Institute of Technology GmbH (ARC)

Donau-City-Straße 1, 1220 Vienna, Austria
Contact: Dr. Denis Havlik, denis.havlik@ait.ac.at

SPACEBEL SA (SPB)

Liège Science Park, B-4031 Angleur, Belgium
Contact: info@spacebel.be

Hochschule für Technik und Wirtschaft des Saarlandes (EIG)

Environmental Informatics Group, Hochschul-Technologie-Zentrum, Altenkesseler Straße 17 Geb. D2, 66115 Saarbrücken, Germany
Contact: sany@enviromatics.net

Open Geospatial Consortium (Europe) Ltd. (OGCE)

8 Coldbath Square, EC1R 5HL, London, United Kingdom
Contact: Martin Klopfer, mklopfer@opengeospatial.org

Fraunhofer Institute IITB (IITB)

Fraunhoferstr. 1, 76131 Karlsruhe, Germany
Contact: Dr. Kym Watson, kym.watson@iitb.fraunhofer.de

MarineTech South Ltd (MTS)

Waterfront Campus, European Way, Southampton, SO14 3ZH, United Kingdom
Contact: Jonathan Williams, jonathan@marinetech.co.uk

IT Innovation Centre (IT-INNOV)

University of Southampton, 2 Venture Road, Southampton, SO16 7NP, United Kingdom
Contact: Dr Z. A. Sabour, zas@it-innovation.soton.ac.uk

BMT Cordah Ltd. (BMT)

Grove House, 7 Ocean Way, Meridians Cross, Ocean Village, Southampton, SO14 3TJ, United Kingdom
Contact: Matthew Rymell, matthew.rymell@bmtcordah.com

Iséo (ISEO)

Technopole Izarbel, 64210 Bidart, France
Contact: Alexandre da Costa, alexandre.da-costa@iseo.fr

KTT-IMA Sarl (KTT-IMA)

20, impasse des Fauvettes, 57460 Behren lès Forbach, France
Contact: Wolfgang Kunz, wkunz_sany@ktt-ima.com

Umweltbundesamt GmbH (UBA)

Spittelauer Lände 5, 1090 Wien, Austria
Contact: Johann Weigl, johann.weigl@umweltbundesamt.at

Soldata SAS (SOL)

Parc de l'île, 21 rue du port 92022, Nanterre cedex 01 41 44 85 00, France
Contact: Sophie Costes, sophie.costes@soldata.fr

Maritime Office in Gdynia (MOG)

81-338 Gdynia, ul. Chrzanowskiego 10, Poland
Contact: Andrzej Królikowski, dum@umgdy.gov.pl

Institut Géographique National (IGN)

73 avenue de Paris, 94165 Saint Mandé Cedex, France
Contact: thierry.person@ign.fr

Gestió d'Infraestructures S.A.U. (GISA)

Carrer dels Vergós, 36-42, 08017 Barcelona, Spain
Contact: Henning Schwarz, hsc@gisa.cat

Institute of Cartography, ETH Zurich (ETHZ)

Wolfgang-Pauli-Str. 15, CH-8093 Zürich, Switzerland
Contact: Prof. Dr. Lorenz Hurni, hurni@karto.baug.ethz.ch