

Live Geography

Standardised Geo-sensor Networks for Real-time Monitoring in Urban Environments

DISSERTATION

for the attainment of the degree
Doktor rerum naturalium

Prepared at the
Geography and Geology
Department at
Paris-Lodron University Salzburg – Faculty of Natural Sciences

In cooperation with
Research Studios Austria and **Massachusetts Institute
of Technology**
Studio iSPACE SENSEable City Lab

Submitted by
Bernd Resch

Reviewed by: Prof. Dr. Josef Strobl
Prof. Dr. Euro Beinat
Prof. Dr. Thomas Blaschke

Salzburg, October 2009

Who believes to be someone, has ceased to become someone.

(Socrates)

*Change will not come if we wait for some other person
or if we wait for some other time.*

We are the ones we've been waiting for.

We are the change that we seek.

(Barack Obama)

Acknowledgements

At first and foremost, I would like to express great gratitude to my family. My parents, my sister and my grandparents were always extraordinarily supportive and encouraging during my entire life in the decisions I took and the paths I chose. This gave me the necessary backing in uncertain times and exhausting periods of re-orientation.

I would like to express special thanks to Manfred Mittlböck from the Research Studio iSPACE, whose broad and far-reaching ideas were of great help and significant importance towards the development of this thesis, conceptually and technically.

Particular thanks is dedicated Josef Strobl for being a helpful, far-sighted and careful supervisor throughout the creation of this thesis. He provided me with comprehensive and precise comments on my research. I would also like to thank Euro Beinat and Thomas Blaschke for being accurate reviewers and proficient co-supervisors.

Then, I would like to thank the Research Studios Austria for offering me the possibility of a well-supported doctorate. A special "thank you" goes to the members of the studio iSPACE, particularly to Manfred John, Michael Lippautz and Martin Dürr for their great work, each of which contributing to technical set-ups within this thesis.

Furthermore, I dedicate my thanks to the MIT SENSEable City Lab for influencing and developing my research on a broad scale. I would particularly like to thank Rex Britter, Fabien Girardin and Eugenio "Huge Lello" Morello. Connected to this, I

would like to express my thanks to Sam Lipson at the City of Cambridge Public Health Department for offering his practical thoughts on urban sensing.

A very special “thank you” is devoted to Jon Reades at UCL, not only for proof-reading this thesis, but also for his comprehensive and far-sighted thematic comments. No matter what, Jon always found the time for offering criticism helping me a lot in the structural and technical development of the thesis. The same applies to Wolfgang John from Chalmers University in Gothenburg. Thanks guys!

Finally, I would like to thank my friends for being supportive and understanding throughout the development of this thesis, and for offering me their time and caring attention. The exhilarating moments with you were an essential motivator throughout the elaboration of the doctorate. I truly appreciate your devotion!

A big “thanks” to all of you !!!

Preface

Sensors have become a vital part of everyday lives. We are surrounded by a great variety of small sensing devices, hidden in our cars, homes, surrounding energy infrastructures, cell phones, and so forth. These sensors extend our human senses and assist us in perceiving our environment in a more holistic way by revealing “unseen” information layers.

But how can we leverage this kind of data? These invisible sensed data have to be made perceivable to the eye – or other human senses. However, creating high quality information layers from raw sensor data is not trivial. User-specific needs in terms of visualisation, data filtering, context presentation, and the way or the medium to convey information have to be carefully assessed.

Moreover, my experiences showed that an increasing number of sensor networks are emerging, but their data are not combinable due to monolithic and closed measurement systems. This prevents trans-domain usage of sensor data and thus calls for new paradigms in terms of data accessibility, information provision and policy-making for real-time data availability.

This thesis is dedicated to tackle that issue from a technical viewpoint. After the beginning of this thesis, also a few large-scale international research projects originated from similar needs. The cooperation with these projects was a valuable help for defining the focus of this thesis, but also a productive interplay by contribution of ideas in both directions.

As this thesis has been conceived over a period of several years, significant pieces of the research outcomes have already been published in a variety of scientific journals and conference proceedings. These papers are listed below. It shall be noted that some sub-parts of the thesis have been directly transferred into scientific papers. Small fractions have been inherited without modification to ensure optimum readability. Mainly, the first three papers of the list below have originated from thesis chapters 1 (Introduction), 6 (“Live Geography” Approach), 7 (Implementation of the Live Geography Approach) and 9 (Discussion and Conclusion). The other papers have emerged from different smaller parts of this thesis integrating ideas, concepts and short text passages.

- Resch, B.**, Mittlboeck, M., Girardin, F., Britter, R. and Ratti, C. (in press) Live Geography – Embedded Sensing for Standardised Urban Environmental Monitoring. *International Journal on Advances in Systems and Measurements*, 2(2), pp. pending.
- Resch, B.**, Mittlboeck, M., Lipson, S., Welsh, M., Bers, J., Britter, R. and Ratti, C. (2009) Urban Sensing Revisited – Common Scents: Towards Standardised Geo-sensor Networks for Public Health Monitoring in the City. In: *Proceedings of the 11th International Conference on Computers in Urban Planning and Urban Management – CUPUM2009*, Hong Kong, 16-18 June 2009.
- Resch, B.**, Mittlboeck, M., Girardin, F., Britter, R. and Ratti, C. (2009) Real-time Geo-awareness - Sensor Data Integration for Environmental Monitoring in the City. In: *Proceedings of the IEEE International Conference on Advanced Geographic Information Systems & Web Services – GEOWS2009*, Cancun, Mexico, 1-7 February 2009, pp. 92-97.
- Resch, B.**, Calabrese, F., Ratti, C. and Biderman, A. (2008) An Approach Towards a Real-time Data Exchange Platform System Architecture. In: *Proceedings of the 6th Annual IEEE International Conference on Pervasive Computing and Communications*, Hong Kong, 17-21 March 2008.
- Resch, B.** and Mittlboeck, M. (under review) Live Geography – Interoperable Geo-Sensor Webs Enabling Portability in Monitoring Applications. In: *Proceedings of the 2nd International Conference on Advanced Geographic Information Systems, Applications, and Services – GEOProcessing 2010*, St. Maarten, Netherlands Antilles, 10-15 February 2010, pp. pending.
- Resch, B.** (2007) Standardised Real-time Information Systems for Environmental Monitoring - An Elaborate Challenge. In: *Proceedings of the GI Days Conference*, Muenster, Germany, 10-12 September 2007.
- Resch, B.**, Schmidt, D. and Blaschke, T. (2007) Enabling Geographic Situational Awareness in Emergency Management. In: *Proceedings of the 2nd Geospatial Integration for Public Safety Conference*, New Orleans, Louisiana, USA, 15-17 April 2007.

- Resch, B.** and Mittlboeck, M. (2007) Einbindung von Echtzeitmessinformation in Geographische Informationssysteme. In: Strobl, J., Blaschke, T., Griesebner, G. (Eds.) (2007) Angewandte Geoinformatik 2007, Wichmann Verlag, Heidelberg, pp. 615-620.
- Mittlboeck, M. and **Resch, B.** (2008) Federal Pervasive Sensor Networks Serving Geographic Information Services. In: Proceedings of the 5th International Symposium on LBS and Telecartography, 26-28 November 2008, Salzburg, Austria.
- Mittlboeck, M. and **Resch, B.** (2009) Live Geography – Real-time Analyse mit GI-Systemen. GISBusiness, 5/2009, pp. 36-39.
- Blaschke, T., Zeil, P., Strobl, J., Lang, S., Tiede, D., Möller, M., Triebnig, G., Schiller, C., Mittlböck, M. and **Resch, B.** (2007) GMES: from Research Projects to Operational Environmental Monitoring Services. International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, XXXVI-1/W51.
- Mittlböck, M., **Resch, B.**, Lippautz, M., Dürr, M. and Kohlmann, M. (2009) Standardisierte Embedded Sensormessnetze zur Integration Lokationsbezogener "Live" Information in GI-Systeme. In: Strobl, J., Blaschke, T., Griesebner, G. (Eds.) (2009) Angewandte Geoinformatik 2009, Wichmann Verlag, Heidelberg, pp. 178-183.
- Mittlboeck, M. and **Resch, B.** (2008) Standardisierte Integration von Real-time Sensormessungen für Zeitnahe GIS-Analyse. In: Strobl, J., Blaschke, T., Griesebner, G. (Eds.) (2008) Angewandte Geoinformatik 2008, Wichmann Verlag, Heidelberg, pp. 112-117.

Affidavit

Herewith I, Bernd Resch, declare that I have written the dissertation at hand fully on my own and that I have not used any other sources apart from those given.

Date

Signature

Key Information

Name:	Bernd Resch
Institution:	Paris-Lodron University of Salzburg, Austria
Department:	Geography and Geology – Geoinformatics
Title:	Live Geography – Standardised Geo-sensor Networks for Real-time Monitoring in Urban Environments
Supervisor:	Prof. Dr. Josef Strobl (Paris-Lodron University of Salzburg)
Reviewers:	Prof. Dr. Euro Beinat (Vrije Universiteit Amsterdam) Prof. Dr. Thomas Blaschke (Paris-Lodron University of Salzburg)
Cooperation Partners:	Research Studios Austria – iSPACE Massachusetts Institute of Technology – SENSEable City Lab

Keywords

1. Keyword: Live Geography
2. Keyword: Urban Monitoring Networks
3. Keyword: Pervasive Geo-sensor Webs
4. Keyword: Standardised Measurement Infrastructure
5. Keyword: Web-based Geospatial Data Analysis
6. Keyword: Real-time Sensor Fusion

Abstract

Real-time monitoring of urban dynamics, which is widely unexplored, has recently received a lot of attention due to the fast rise of inexpensive pervasive sensor technologies. Following the comprehensive vision of a “digital skin for planet Earth” by Neil Gross (1999), it can be presumed that geo-sensor web deployments will experience a considerable boost within the coming years, as pervasive sensing has recently become feasible and affordable. This enriches knowledge about our environment with uncharted real-time information layers.

However, leveraging sensor data in an ad-hoc fashion is not trivial, as ubiquitous geo-sensor web applications comprise numerous technologies, such as sensors, communications, massive data manipulation and analysis, data fusion with mathematical modelling, the production of outputs on a variety of scales, the provision of information as both hard data and user-sensitive visualisation, together with appropriate delivery structures. Furthermore, current sensor networks are often built up in monolithic and closed “black box” systems. Thus, measurement data are mostly not inter-combinable because of lacking interoperability and deficient coordination.

This thesis focuses on pervasive sensing in urban environments, which poses very particular challenges – as well technical and technological ones, as socio-political and privacy-related ones. The *Live Geography* approach presented in this thesis seeks to tackle these challenges with an open sensing infrastructure for urban monitoring applications. The system makes extensive use of open (geospatial) standards throughout the entire process chain – from sensor fusion to data analysis, Complex Event Processing (CEP), alerting, and finally visualisation. The thesis discusses the implemented modules as well as the created architecture as a whole.

Finally, it is shown how the methodology can potentially influence the city and its inhabitants by “making the abstract real”. In other words, the thesis illustrates how pervasive monitoring infrastructures can change urban social interactions and people’s short-term behaviour, and which issues are related to establishing such systems.

Thesis Layout




Table of Contents

1	Introduction.....	1
1.1	Augmenting Research on Multi-layered City Dynamics	2
1.2	Pervasive Urban Monitoring – The Technical View	3
1.3	Closed Monolithic Systems Limit Monitoring Capacities	4
1.4	Standardisation Enables Open Measurement Infrastructures.....	5
1.5	Summary and Approach	6
2	Thesis Objectives	9
2.1	Research Question	9
2.2	Goals for the Live Geography Infrastructure	10
2.3	Goals for the Implementation of the Live Geography Approach	12
2.4	Non-Goals.....	14
3	Thematic Context of the Thesis	17
3.1	Live Geography – Disambiguation.....	17
3.2	The “City” as a Functional Context.....	19
3.3	Particular Challenges of Urban Monitoring	22
3.3.1	Technological and Technical Issues	22
3.3.2	Various Stakeholders	23
3.3.3	Impact and Privacy Implications	24
3.4	Sensor Web Technology	25
3.5	Why Standards?.....	27
3.5.1	Standardisation Initiatives	27
3.5.2	Benefits Arising from Usage of Standards.....	28
3.6	Policy-frameworks for the Integration of Real-time Sensor Information ..	29
3.7	Use Cases and Application Scenarios.....	31
3.7.1	Use Cases	31
3.7.2	Application Scenarios	32
4	Related Work	42
4.1	Sensor Networks and GIS.....	42
4.2	Open Sensor Data Integration Architectures.....	43
4.3	Empirical Aspects – A State of the Art Review	44
4.3.1	Limitations	44
4.3.2	Technical Deliberations	44
4.3.3	External Dependencies.....	45
4.3.4	Security Aspects	45

5	Methods	51
5.1	Urban Geo-Sensor Networks	51
5.1.1	Literature Review	51
5.1.2	Sensor Web Characteristics.....	53
5.2	Real-time Data Integration and Sensor Fusion for GIS	57
5.2.1	Literature Review	57
5.2.2	Shortcomings of Previous Approaches	58
5.3	Web-based Geospatial Data Analysis.....	60
5.3.1	The OGC WPS Standard	60
5.3.2	Literature Review – WPS Implementations	60
5.3.3	WPS Profiles	61
5.3.4	Asynchronous processing	62
6	“Live Geography” Approach	65
6.1	Background	65
6.2	Standardised System Architecture	67
6.2.1	Web Service Standards	68
6.2.2	Sensor Web Enablement Initiative	69
6.2.3	Other Relevant OGC Standards	76
6.3	Geo-workflow Disambiguation – Sensor Fusion, Data Fusion and Geo-processing	83
6.3.1	Background – Literature.....	83
6.3.2	Data Fusion Definitions	85
6.3.3	Geo-processing Definitions	86
6.4	Geo-sensor Webs	90
6.5	Real-time Data Integration – Sensor Fusion	91
6.6	Location-aware Event-based Alerting	93
6.7	Real-time Geo-processing Using Server-based Analysis Technology	97
7	Implementation of the Live Geography Approach	101
7.1	Embedded Sensing Device.....	101
7.1.1	Sensor Pod Hardware.....	103
7.1.2	Software Infrastructure	105
7.1.3	Embedded Database	108
7.2	Standardised Service Infrastructure	111
7.2.1	Embedded OGC Sensor Observation Service.....	111
7.2.2	Embedded OGC Sensor Alert Service.....	113
7.3	Real-time Data Integration – Sensor Fusion	114
7.3.1	Technical Realisation.....	114
7.3.2	Functionality.....	114

7.3.3	Innovation and Benefits	117
7.4	CEP and Location-aware Complex Event Processing and Alerting	118
7.4.1	CEP Engines	118
7.4.2	Geographical Code Libraries.....	119
7.4.3	Technical Realisation.....	120
7.5	Live Geospatial Data Analysis.....	123
7.5.1	Desktop GIS Processing of Sensor Information	123
7.5.2	Time-series Analysis	124
7.5.3	Web-based Geospatial Processing	125
8	Results	129
8.1	Results Summary	130
8.2	The Urban Environment as the Functional Context	131
8.3	Live Geography Approach.....	132
8.4	Benefits of the Developed Approach	133
8.5	Implementation of the Live Geography Approach	135
8.5.1	Embedded Geo-Sensing Device	135
8.5.2	Standardised Service Infrastructure	135
8.5.3	Real-time Data Integration – Sensor Fusion	136
8.5.4	Complex Event Processing and Alerting	136
8.5.5	Web-based Data Analysis	137
8.6	Live Geography Portability – Implemented End Applications	137
8.6.1	Live Pollutant Monitoring for Public Health.....	138
8.6.2	Fine-grained Air Temperature Variations.....	140
8.6.3	Ubiquitous Biometric Parameter Surveillance.....	141
8.6.4	Real-time Air Quality Assessment	142
9	Discussion and Conclusion	146
9.1	Future Outlook – Complete Live Geography Infrastructure.....	147
9.1.1	Further Research Directions	147
9.1.2	Fully-fledged Live Geography Infrastructure	149
9.2	Effects on the Urban Environment.....	151
9.3	Changing the Perception of the City	154
9.4	Conclusion	157
	References.....	161
	List of Abbreviations.....	177
	List of Figures.....	180
	List of Tables	182
	Annex I – XML Structures	183

INTRODUCTION

INTRODUCTION

1 Introduction

"In the next century, planet earth will don an electronic skin. It will use the Internet as a scaffold to support and transmit its sensations. This skin is already being stitched together. It consists of millions of embedded electronic measuring devices: thermostats, pressure gauges, pollution detectors, cameras, microphones, glucose sensors, EKGs, electroencephalographs. These will probe and monitor cities and endangered species, the atmosphere, our ships, highways and fleets of trucks, our conversations, our bodies – even our dreams." (Gross, 1999)

Following this comprehensive vision by Neil Gross (1999), it can be assumed that sensor network deployments will increase dramatically within the coming years, as pervasive sensing has recently become feasible and affordable. This enriches knowledge about our environment with previously uncharted real-time information layers.

However, leveraging sensor data in an ad-hoc fashion is not trivial as ubiquitous geo-sensor web applications comprise numerous technologies, such as sensors, communications, massive data manipulation and analysis, data fusion with mathematical modelling, the production of outputs on a variety of scales, the provision of information as both hard data and user-sensitive visualisation, together with appropriate delivery structures. Apart from this, requirements for geo-sensor webs are highly heterogeneous depending on the functional context.

This thesis focuses on pervasive sensing in urban environments, which poses very particular challenges – as well technical and technological ones, as socio-political and privacy-related ones. This distinguishes urban monitoring applications clearly from other, less central scenarios like for instance ice thickness monitoring in the Antarctica (Delin et al., 2003) or gravitational mass-movement tracking for change detection of alpine permafrost (Paulsen, 2008).

Summarising, the main question to answer is: how can interoperable infrastructures, which flexibly integrate geo-sensor networks, real-time sensor fusion and geospatial data processing, influence city dynamics by fostering the deployment of monitoring applications in urban environments and by complementing or even replacing monolithic measurement systems?

To answer this question, the technological prerequisites for ubiquitous monitoring have to be elucidated and examined in the context of urban environments. Thus,

the first step to set up the frame for this thesis is to investigate cities and their particular multi-layered functionality.

1.1 Augmenting Research on Multi-layered City Dynamics

Projections stated that the wide-spread distribution of high-speed internet connections will render geographical distance irrelevant (Cairncross, 1997 and Borsch, 2009), and that cities are not more than mere artefacts of the industrial age (Gilder and Peters 1995). As a side effect, cities were presumed to drastically decrease in importance as physical and social connections, and would play an increasingly ancillary role in socio-technical research.

In reality, the world developed completely differently – cities are back. In fact, a United Nations (UN) report, which has been released before the World Population Day in 2007, states that for the first time in history, more people now live in cities than rural areas (United Nations Population Fund, 2007). Thus, cities in their multi-layered complexity in terms of social interactions, living space provision, infrastructure development and other crucial human factors of everyday life have re-gained importance in scientific research. This arises from the fact – amongst others – that major developments of scientific and technological innovation took place in the urban context (Dierig et al., 2000 and Netherlands Organization for Scientific Research, 2007).

Here, especially real-time monitoring of urban processes is widely unexplored and has recently received a lot of attention due to the fast rise of inexpensive pervasive sensor technologies, which made ubiquitous sensing feasible and enriches research on cities with uncharted up-to-date information layers.

The reason for this lack of experience is that ambient and continuous monitoring is an enormous challenge, and this is particularly true in the urban context, which poses very specific challenges. These comprise well-known technological questions, but also significant social and political ones, which are rapidly gaining importance.

Figure 1 illustrates the monitoring concept of a *WikiCity* (SENSEable City Lab, 2009), which interprets the city as a “control system”, where the objects connect the physical environment with the digital space.




Figure 1: Urban Monitoring – the Connection between the Physical and the Digital World.
(SENSEable City Lab, 2009)

This connection between the physical and the digital worlds illustrated in Figure 1 can be established by ubiquitous sensor networks creating a digital representation of the real world. Consequently, these sensors introduce the possibility of a feedback loop – because sensor systems do not only view and analyse the city as a functional entity, but they can also influence urban processes in real-time, by coupling them with decision support mechanisms, location-based services (LBS) or other spatio-temporal technologies.

1.2 Pervasive Urban Monitoring – The Technical View

From a technical viewpoint, pervasive sensing encompasses a wide range of diverse technologies. As mentioned in the introductory paragraphs, this involves sensors, communication infrastructure, data manipulation and analysis, sensor fusion, the creation of meaningful outputs, the provision of information in user-sensitive visualisations, and appropriate delivery structures. The present thesis addresses the nature of this supply chain by proposing a fully standardised infrastructure.

Additionally, a key requirement for achieving meaningful analysis results in geo-processing systems is high quality of underlying data. Quality criteria for geospatial data comprise accuracy, completeness and topicality. While the first two parameters have traditionally received a lot of attention, the topicality parameter is still neglected in most cases. This is partly due to the fact that geo-data is per definition historic and so far, geospatial processing has focused on analysing static data, with low temporal fluctuations. (Resch et al., 2009a)

Furthermore, an overarching aspect is that all components of measurement systems are currently undergoing both great performance enhancement combined with drastic price reduction (Paulsen and Riegger, 2006). This has led to the deployment of a number of geo-sensor networks, as described in sub-chapter 5.1.

On the positive side the growing establishment of such networks will further decrease prices and improve component performance. This will particularly be so if the monitoring policies enforced by super-national directives move from a mathematical modelling base to a more pervasive monitoring structure. Thus, it can be assumed that sensor network deployments will increase dramatically within the coming decade – especially in cities.

This raises the issue that real-time integration of heterogeneous data sources will be of particular importance in urban environments as a broad spectrum of diverse actors establishes their systems, often measuring the same parameters. This evidently distinguishes urban monitoring applications clearly from other, less connected scenarios such as ice thickness monitoring in the Antarctica (Delin et al., 2003) or gravitational mass-movement tracking for change detection of alpine permafrost (Paulsen, 2008).

1.3 Closed Monolithic Systems Limit Monitoring Capacities

Consequently, ubiquitous monitoring is a critical process in cities to ensure public safety including the state of the national infrastructure, to set up continuous information services, and to provide input for spatial decision support systems.

However, establishing an overarching monitoring infrastructure is not trivial. Currently, different authorities with heterogeneous interests each implement their own monolithic systems to achieve very specific goals. For instance, regional governments measure water levels for flood water prediction, while local governments monitor air quality to dynamically adapt traffic conditions, and energy providers assess water flow in order to estimate energy potentials.

As shown in Figure 2, these measurement systems are built up in a minimalistic – though often short-term cost-optimised – manner resulting in a hard-wired workflow. Such systems mostly serve a single purpose and can often not be adapted to other end applications. This also means that the components within the overall process are chained together in a fixed way (“black box”), which prevents re-use of single modules.




Figure 2: Current Approach – Monolithic Single-purpose Sensor Systems.

These single-purpose measurement systems as shown in Figure 1 are typically characterised by a centralised sensor network sending data in a proprietary format, which prevents distributed architectures and interoperability with other systems. The gained measurements are stored in a database, which can often become a bottle-neck in the overall workflow when dealing with large amounts of data.

Before the sensor data can be analysed they have to be pre-processed. However, this pre-processing step has to follow strict and rigid rules in order to fit the specific non-recurring interfaces of the analysis system. Such analysis systems mostly process data in a closed “black box” model. Furthermore, they usually provide data in a singular and application-tailored format preventing open use and re-use of analysed data.

The fact that these systems tend to be deployed in an isolated and uncoordinated way means that the automatic assembly and analysis of these diverse data streams is impossible. However, making use of all available data sources is a prerequisite for holistic and successful monitoring for broad decision support in an urban context. This applies to emergency situations as well as to continuous observation of environmental parameters.

1.4 Standardisation Enables Open Measurement Infrastructures

The widespread availability of sensor data with high spatial and temporal resolution is expected to increase dramatically with rapidly decreasing prices (Resch et al., in press), particularly if costs are driven down by mass utilisation. One way to address this issue is the extensive use of open standards and geospatial web services for structuring and managing heterogeneous data. In this context, the main challenge is distributed processing of vast amounts of sensor data in real-time.

From a political and legal standpoint, national and international legislative bodies are called upon to foster the introduction of open standards in public institutions. Comprehensive early efforts in this direction have been made by the European Commission (EC) through targeted directives (s. sub-chapter 3.6). These include

the INSPIRE (INfrastructure for SPatial InfoRmation in Europe) initiative, which aims at Europe-wide harmonisation of discovery and usage of geographical data for analysing and solving environmental problems (Commission of the European Communities, 2009a).

These regulations support the establishment of ubiquitous and generically applicable real-time data integration mechanisms. Shifting development away from proprietary single-purpose implementations towards interoperable analysis systems will not only enable live assessment of the urban environment, but can also lead to a new perception of the city by its inhabitants. Consequently, this trend may in turn foster the creation of innovative applications that treat the city as an interactive sensing platform. One example is the *WikiCity* concept (Resch et al., 2008), involving the people themselves into re-shaping the urban context.

1.5 Summary and Approach

Of specific interest in this dissertation is the concern that most sensor networks are being built up in monolithic and specialised application-centred measurement systems. Subsequently, there is a clear gap between sensor network research and largely heterogeneous end user requirements. Sensor network research is usually dedicated to a long-term vision, which tells a compelling story about potential applications. On the contrary, the actual implementation is often not more than a very specific use case oriented closed system without taking into account technical and conceptual issues such as interoperability, sustainable development, portability, or combination with established data analysis systems.

Therefore, the availability of geo-sensor networks is growing but still limited. Deborah Estrin pointed out in 2004 that no real sensor network applications exist, apart from short-lived prototypical and very domain-specific demos (Xu, 2004). Recently, some examples of urban geo-sensor networks arose (s. sub-chapter 5.1), but the clear social, health and economic benefits have not been demonstrated or described in a manner that would compel this sort of investment in particular for urban environments.

The goal of this thesis is that the resultant highly flexible architecture will bring sensor network applications one step further towards the realisation of the vision of a “digital skin for planet earth” (Gross, 1999). In other words, the broad objective of the thesis is to develop an overarching infrastructure for various kinds of sensor network applications in urban environments. The developed architecture can potentially have far-reaching impacts on urban monitoring systems through the deployment of ubiquitous and fine-grained sensor networks.

Chapter Summary

Pervasive urban monitoring comprises sensors, communications, massive data manipulation and analysis, data fusion with mathematical modelling, the production of outputs on a variety of scales, the provision of information as both hard data and user-sensitive visualisation, together with appropriate delivery structures.

However, setting up an overarching monitoring system is not trivial. Currently, different authorities with heterogeneous interests each implement their own monolithic infrastructures to achieve specific goals. Furthermore, the topicality parameter (apart from accuracy and completeness) has only recently gained importance in geospatial analysis through the emergence of a variety of real-time data sources such as sensor networks or RFID-based systems, which make pervasive monitoring networks feasible and affordable.

Thus, continuous monitoring is an enormous challenge, and this is particularly true in the urban context, which poses very specific challenges. These comprise well-known technological issues, but also significant social and political ones, which are rapidly gaining importance.

This thesis presents the Live Geography approach, which proposes a fully standards-based distributed infrastructure combining current sensor data and server-based analysis systems for monitoring applications in urban environments. Its objective is to answer the question how interoperable infrastructures, which flexibly integrate geo-sensor networks, real-time sensor fusion and geospatial data processing, can influence city dynamics by fostering the deployment of monitoring applications in urban environments and by complementing or even replacing monolithic measurement systems.

THESIS OBJECTIVES

THESE OBJECTIVES

2 Thesis Objectives

The main aim of this thesis is the development of an approach for a standardised sensing and data analysis infrastructure targeted to monitoring applications using distributed geo-sensor networks, with a particular focus on urban environments. The architectural design has to be highly interoperable, modular and flexible, so that several possible real-time geo-analysis applications integrating sensor data can be implemented upon it. Furthermore, a sample application is implemented for demonstration purposes. Apart from these technical outcomes, possible impacts on urban dynamics and processes will be briefly discussed.

The following sub-chapters comprise a description of the essential research questions as well as the differentiated elucidation of the project goals and non-goals for the proposed monitoring infrastructure and the sample implementation.

2.1 Research Question

As mentioned in chapter 1, current sensor networks are mostly built up in monolithic and closed systems. They often lack interoperability and open accessibility, thus limiting trans-domain usage of gathered data in analysis workflows. The overarching high-level research question of this dissertation arising from this issue is:

How can interoperable infrastructures, which flexibly integrate geo-sensor networks, real-time sensor fusion and geospatial data processing, influence city dynamics by fostering the deployment of monitoring applications in urban environments and by complementing or even replacing monolithic measurement systems?

This multi-dimensional question can be sub-divided into several strategic questions and research task oriented operational questions:

Strategic Questions

- Why is interoperability an essential requirement for monitoring architectures and how can it be accomplished for geo-sensor networks?
- Which modules does a flexible geo-sensor data analysis workflow have to contain?

- How can the functionality of the infrastructure be expressed in terms of a workflow chaining together different service interfaces using (geospatial) standards?
- How can real-time data integration be achieved for geo-workflows?
- Which implications does “pervasive urban monitoring” contain for the individual?
- What societal, technological and perceptual impacts do pervasive monitoring infrastructures have on the urban environment?

Operational Questions

- What are monolithic measurement infrastructures and how do they limit sensor network development and deployment?
- What does the term “city” mean in terms of a functional application context for monitoring applications?
- What are particular challenges of urban monitoring?
- Which geospatial standards have to be taken into account for the creation of an interoperable measurement and data analysis infrastructure?
- What are the benefits of open interoperable sensing architectures?
- Which effects on the city and its inhabitants can potentially be expected using such systems?

2.2 Goals for the Live Geography Infrastructure

The overarching goal of the dissertation is to conceive a widely **interoperable infrastructure** for distributed and decentralised geo-sensor web applications with a special focus on urban environments. In this regard, a central aspect is the design of a highly modular architecture, which makes the single elements in the workflow inter-chainable. Therefore, this component-oriented approach allows for flexible module combination and their integration into existing geo-data analysis workflows.

To achieve this component orchestration, the architecture design requires provision for numerous **international standards** to guarantee far reaching compliance and applicability in a possibly world-wide extent. This need arises from the fact that sensor web applications might be based on very different sensor models, satellite types or communication technologies, which have to be supported. Moreover, the technological basis has to be open for scientific advances in sensor technologies,

measurement mechanisms, and computation power to ensure extensibility and scalability.

Furthermore, an essential project goal is widespread compatibility of the basic system with well-known and ubiquitously used **web service standards** (W3C internet standards). This compliance allows for superimposing numerous applications upon the developed general infrastructure to use sensor webs in connection with web services. This is a highly functional and useful combination as retrieved or rehashed data can be made publicly available to a broad user group.

Another central goal is the creation of a mechanism to enable real-time integration of heterogeneous data sources. The developed infrastructure has to contain a dedicated component for **sensor fusion** and data integration to harmonise a wide range of different data sources including in-situ sensors, archived measurements and external non-real-time data sources.

The two latter goals mentioned above play a significant role towards the aim of providing for the possibility of **integrating existing sensor networks**. These legacy measurement systems are often conceived to serve a single purpose and are thus not designed for portability and open accessibility. Despite these shortcomings of existing sensor networks, the infrastructure developed in this thesis has to be flexible enough to integrate their data in order to achieve maximum data availability for further analysis.

This ambition is strongly coupled with the objective of creating a defined mechanism for **sensor data provision**. Arranging for uniform data representation facilitates the integration of sensor measurements into a wide variety of applications for data analysis, visualisation or storage. Data provision will be realised by the use of well-established service interfaces for different geospatial data types such as geographic features, image maps or geographic coverages.

Regarding geo-sensor network design and deployment issues, **particular challenges of establishing monitoring systems in urban environments** have to be investigated apart from general sensor web research subjects. It is necessary to elucidate not only well-known technical questions, but also other challenges, which appear particularly in the environment of the city. These comprise social, political and privacy issues, or other problems arising from the involvement of heterogeneous interests.

Naturally, a basic **security concept** has to be created in order to protect transmitted data from being spied on. However, higher level security such as the actual implementation of intra-sensor communication encryption and secure data

transport to a base station or central processing entity is not a matter of this generalised basis. This issue shall be individually solved for each application, as mentioned in sub-chapter 2.4.

2.3 Goals for the Implementation of the Live Geography Approach

The sample implementation of the proposed infrastructure will comprise different software and hardware modules, ranging from a sensor pod to real-time integration, event processing, alerting, and data analysis components. The prototype shall dedicatedly support mobile and stationary sensors, and will contain a sample hardware configuration and software implementation. The key deliberations concerning this practical realisation will be elucidated in this sub-chapter.

In addition to the goals for the open measurement architecture described in sub-chapter 2.2, several practical considerations have to be made before the realisation of a sample application.

First, **integrative processing of complementary sensor data** sets has to be allowed for. This means that the data exchange structures and methods must not be dependent on certain types or models of sensors. Moreover, general sensor web criteria have to be fulfilled to guarantee wide applicability and compatibility with other systems. These criteria encompass interoperability, scalability and intelligence, as illustrated in sub-chapter 5.1.

The **sensor device** itself has to comprise a central processor, a communication unit and a small memory of about 16MB, which is the minimum configuration for the installation of an embedded operating system. Furthermore, a small-footprint web server and a SQL-conformant database have to be contained. Moreover, the device has to feature several interfaces (serial, USB, I²C etc.) for different sensors and an embedded Global Positioning System (GPS) module, which is of special importance for mobile sensors. These features allow for storing a measurement history, which is necessary for on-site quality assurance, error detection, and short-term trend analysis.

Considering the aspect of **live data provision and analysis**, near real-time requirements have to be met in the sense of “as fast as possible” with a maximum delay of a few (tens of) seconds for simple data analysis algorithms – depending on the particular use case (cp. sub-chapter 3.1).

Tightly connected to the real-time problem is the question, **which data should actually be sent by the sensor**. For the Live Geography implementation, as well

mobile as stationary sensors should always include their geographical positions when transmitting sensor measurements. Also, sensors have to support filtering mechanisms such as spatial, temporal or measurement property filters.

A challenge related to the latter requirement is live data integration. In order to create up-to-date information layers from raw sensor data, different data sources have to be harmonised in near real-time. Thus, a **sensor fusion module** is a central goal of the implementation as it is the main component connecting the sensor network with data analysis mechanisms. It has to integrate measurements in real-time and output harmonised data in a well-established and standardised data format.

At the other end of the workflow, several **analysis and visualisation components** will be implemented. These are responsible for illustrating complex issues in a simple and intuitive manner. In other words, the visualisation modules display outputs of analysis processes, which combine various kinds of input data in pre-defined algorithms into a single user-tailored output. These visualisations can for instance serve as a basis for spatial decision support.

Naturally, security questions have to be addressed when dealing with sensor networks and geospatial data analysis. However, this is not a central research topic of this thesis. As a variety of **security mechanisms** already exists, only a simple proof-of-concept will be implemented comprising secure communication between the sensor network and a base station, e.g. a sensor fusion server. Essential security facets will be discussed in sub-chapter 4.3.

Next, the pod's **hardware design** is an important issue. It has to be mentioned that at any rate, a trade-off between the pod's size, including e.g. antenna size or transducers, and its longevity and capabilities has to be made. In general, it can be stated that the pod has to be designed according to its needs presupposed by the application scenario. These requirements comprise ease of deployment, robustness, service intervals, life cycle etc. In either case the pod's application logic should include a power save mode so that energy is only consumed during the actual communication and measurement processes, which significantly extends the battery's lifespan.

Obviously, also the **robustness** against mechanical and physical impacts from outside, such as oppressive human attacks or harsh weather conditions, has to be taken into account when producing a sensor pod. This is particularly true when creating a critical highly available and reliable sensor web as described by Baker and Carpenter (2005). Furthermore, the entire system has to operate continuously over a long period of time in a reliable and faultless manner mostly in harsh

environments, e.g. at extreme temperatures, vibration or at high humidity, to preserve meaningful outputs. This presumption is especially valid for mobile sensors facing changing environments, which have to be constructed in a robust way to bear extreme changes of outside conditions. However, these mentioned requirements are highly application-specific and may vary fundamentally. Thus, the implementation presented in this thesis has to take those factors into account, but they are not a strict requirement in the system design.

2.4 Non-Goals

As the infrastructure approach pursued in this thesis is comprehensive in terms of functionality, it requires trade-offs in research areas, which are non-essential for the proposed solution. In other words, the single developed components do not claim to be immensely cutting edge by themselves, but the uniqueness of this thesis is their integration in an interoperable infrastructure. Thus, a clear distinction has to be made from other open sensing architecture approaches (s. chapter 4).

The first definite non-goal is that the implementation of the sensor pod will not have to take into account the **whole spectrum of technical finesses** such as efficient bandwidth utilisation, optimised sensor case design and distributed radio access mechanisms in the network. These requirements vary substantially depending on the specific implementation and can thus not be generically answered.

As mentioned in sub-chapter 2.3, **minimised hardware design** is of crucial importance for pervasive monitoring applications. However, within this thesis, it is not an essential requirement for answering the defined research questions. Thus, it shall be taken into account when establishing the system, but it shall not be a substantial quality factor.

Moreover, **cameras shall not be considered a type of remote sensors**. The usage of cameras in urban environments raises a lot of questions related to privacy, data ownership, data storage and deployment issues, which are not in the scope of this thesis.

Next, no **sophisticated security solution** tailored to application-specific needs will be developed. Security per se is a comprehensive research area in terms of authentication, authorisation, data integrity, encryption technology etc., which is not essential for the goals of this thesis. Thus, the sensor device shall only contain a basic security solution to prove the concept.

Furthermore, no in-depth investigation **of available geospatial processing technologies** will be conducted. The proposed infrastructure shall be generic enough to incorporate various Geographic Information System (GIS) analysis software. For demonstration purposes, a basic processing algorithm (e.g. interpolation) shall be implemented using existing GIS software.

More general design topics when building a sensor web are **environmental issues**, such as calibration over a long time period for long-term predictions, will not be discussed in detail. Naturally, these aspects are particularly important for deployments comprising a big number of sensor pods, which are perhaps abandoned after they are disused. This facet will not be treated separately due to its multi-dimensional complexity and its limited relevance for the goals of this thesis.

Chapter Summary

The central question of this thesis is how interoperable infrastructures, which flexibly integrate geo-sensor networks, real-time sensor fusion and geospatial data processing, can influence city dynamics by fostering the deployment of monitoring applications in urban environments and by complementing or even replacing monolithic measurement systems.

The answer to this question is manifold. It requires the investigation of the functional context of the city, an assessment of particular challenges for monitoring systems in urban environments, the identification of shortcomings of current approaches, the development of a highly flexible and standardised infrastructure for real-time sensor fusion and data processing, and the examination of its implications for urban processes.

However, definite non-goals of this thesis are optimised hardware and software development, in-depth assessment on privacy implications and the establishment of a comprehensive use case tailored security model.

THEMATIC CONTEXT

THEMATIC CONTEXT

3 Thematic Context of the Thesis

This chapter serves for defining the thematic frame of the thesis at hand. It establishes a context for the research questions listed in chapter 2, and thus describes the environment, within which the Live Geography approach is developed.

The chapter comprises a disambiguation of the term “Live Geography”, illustrates how the city represents the functional context including particular challenges of urban monitoring, discusses underlying technological developments (sensor webs and geospatial standards), and finally concludes with a description of use cases and potential application scenarios for the concept developed within this dissertation.

3.1 Live Geography – Disambiguation

This thesis presents the *Live Geography* approach, which proposes a fully standards-based distributed infrastructure combining live measurement data with sensor fusion mechanisms and server-based analysis systems for monitoring applications in urban environments.

Live Geography Origins

The term “Live Geography” originates from the modification of the expression “Real-time Geography” as definitions and usages of the latter expression are vague and vary on a quite broad scale. Fischer (2007) defines Real-time Geography as the process of “collecting spatial information about our living space or about our own activities by ourselves and sharing it with our community in real-time”. Another description of Real-time Geography by Ride the Lobster (2007) is “knowing where you are on planet earth, to the nearest metre and within a fraction of a second, and telling something about you and your environment such as your temperature, heart rate or body movement”.

Apart from these definitions, the term is generally understood as providing spatial information in a timely manner without necessarily accounting for a feedback loop or dynamic processes.

In these interpretations of the expression “real-time”, it has been strongly mitigated. The term “real-time” originated in the field of computer science, where it initially described a process, which was completed “without any delay”. This broad view was then divided into hard and soft real-time demands. Soft real-time basically defines that deadlines are important, but the whole system will still function correctly if deadlines are occasionally missed. The latter is not true for hard real-time systems. Another term to express non-rigorous temporal

requirements is “near real-time”, which describes a delay introduced into real-time applications, e.g. by automated data processing or data transmission (General Services Administration, 1996). Hence, the term accounts for the delay between the occurrence of an event and the subsequent use of the processed data.

The Term “Real-time” in a Geographic Context

However, these definitions of the term “real-time” have been set up for the domain of computer science. Thus, it is important to evaluate and re-define the expression in the context of geography. Naturally, strict real-time requirements are a central aspect in monitoring applications, whereby these demands are highly application-specific and can vary significantly. Therefore, they are not a fundamental goal in the field of geography, as the term “real-time” is primarily defined by an “exact point in time”, which is the same for all sensors to create a significant measurement outcome. Secondly, the term defines the possibility to start a synchronous conversation at a certain time, which might often be important for geographical monitoring applications, e.g. to enable the generation of an exact development graph for temporal pollutant dispersion over a defined period of time in precise intervals.

Additionally to the suggestion of assessability of the environment in the “now”, the expression “Live Geography” also implies a feedback loop. The term “Geography” does not only define the description of location-aware parameters, but also entails the exploration of causal patterns in these data. In the context of geo-sensor network applications, this in turn means that the urban environment is not only analysed remotely by examining quasi-static data, but the procedure of sensing and processing live data offers the possibility of modifying the urban context in an ad-hoc fashion.

Conclusion

In conclusion, it can be stated that the strict term “real-time” can be interpreted as “at present” for urban monitoring applications, in the sense that the aim is to assess the environment “now”, not a historical and perhaps outdated representation. However, these topicality requirements can vary depending on the application context. For instance, an update on traffic conditions does not have to exceed a delay of a couple of minutes when this information is used for navigation instructions, whereas a 30 minute update interval can well be sufficient for short-term trip planning.

To account for this non-rigorous requirement, the term "Live Geography" seems better suited than "Real-time Geography". In this reflection, "near real-time" appears to be closest to „live“, as it does not impose rigid deadlines and the expression itself suggests dynamic adaptation of a time period according to different usage contexts.

3.2 The "City" as a Functional Context

As mentioned at the beginning of the introductory chapter, the term "city" comprises not only a geographical area characterised by a dense accumulation of people or buildings. More, it is a multi-layered construct containing multiple dimensions of social, technological and physical interconnections.

Multi-dimensional Complexity

Thus, as Zardini (2006) states, a variety of terms have emerged to describe the evolving urban environment, including virtual city, city of bits, event city, cyber city, global city, network city, or renewable city. Naturally, these terms depend on the specific viewpoint representing the complexity of the phenomena being observed and analysed.

Peter Hall (1966) tries to characterise this multi-dimensional complexity covering culture, politics, trade, communications infrastructure, finance, technology and universities. In more recent work, Hall and Pfeiffer (2000) state that a liveable city has many facets, which revolve around quality of life. These are living space, elementary infrastructure, traffic and land utilisation.

The City as a Multifaceted Process

Manuel Castells (1996) pursues another approach describing a city to be "not a place but a process". In this context, processes are considered the connections between centres in a global network. More, Castells imagines the city as a spatial system of advanced service activities, and claims that information and communication networks constitute the modern social morphology of our societies in the informational age – as opposed to the industrial age.

Consequently, Castells argues that places of operation are called nodes and hubs. Nodes are places where important functions are performed linking localities into the whole network. On the other hand, communication hubs are places coordinating interactions across the network. Castells concludes that the analysis of (global) cities yields the most direct illustration of the space of flows in nodes and hubs.

A similar but more economy-driven view is presented by Friedmann (1986), who states that cities are the basing points of capital, and the resulting linkages create a complex spatial hierarchy. In his interpretation, the hierarchy is formed by taking into account a number of city characteristics: the importance of the city as a finance centre, corporate headquarters, international institutions, business services, manufacturing, transportation, and population size.

A very energy-centred description of urban environments is done by Droege (2006). The author sees the common aspect of cities in their foundation as “creatures of their energy regimes”. Moreover, Droege asserts that a large part of global financial transactions, trade, command and control, and cultural production occur in and among cities.

A Technology-driven View of the City

Mitchell (1996) also lays out a multi-dimensional definition of the term city, which is strongly motivated by technology. He states that the future city will be unrooted to any definite geographic place on the surface of the earth, constrained by connectivity and bandwidth issues rather than by physical accessibility and land values, widely asynchronous in its operation. The inhabitants are not humans, but agents, which Mitchell describes as “collections of aliases and agents”.

Places in the “City of Bits” will be constituted virtually by pieces of software instead of physically, and they will be connected by logical inter-linkages rather than by doors, passageways and streets. He imagines that computer networks become as elementary to urban life as street systems. According to the author, memory and screen space will become the rare resource, like a sort of future real estate. In consequence of the technological predominance, much of the economic, social, political, and cultural action shifts into cyberspace.

Assessment of Urban Dynamics

Despite their diversity of approach, all of the above studies reveal how the cities, in which we live, have changed. There is not a single universal and common way of describing cities and urbanity, but ways of exploring the city have changed together with their inhabitants. This is one reason why people have become the new centre of interest, as stated by Zardini (2006).

From an investigative viewpoint, there is growing interest in the environmental and ecological issues within cities, especially sustainability and bio-diversity. This is mainly because of a more profound faith in the efficacy of the tools and methods of urban design and monitoring, which can foster the understanding and re-design of

the urban environment (Zardini, 2006). Thus, an upcoming challenge is not barely a re-adjustment of the technical instruments, but rather a conceptual change: a shift happened from Bocconi's "city that rises" to the "landscape that advances" providing more possibilities in monitoring and re-designing a city, than just in building it up.

In this context, it is evident that some of the most innovative approaches of the last decades exceeded existing technical knowledge and competence to realise them at the time. However today, they seem not only relevant, but also feasible as informed real-time decisions transform both the atmosphere of cities and the hierarchy of our priorities on a daily basis. For example, we are already using mobile smart phones, ubiquitous personal information devices, to coordinate and adjust our plans on-the-fly and at a distance by receiving up-to-date information on our environment.

Augmenting Human Senses

This also means that the sensorial dimension beyond the prevalence of the vision – acoustic, olfactory, touch and other senses – which have not yet been integrated into urban spaces, shall play an increasingly important role towards the development of "sensorial city planning". (Zardini, 2006)

This integration of additional senses is not only relevant for city planning, but also for city management. To obtain a complete picture of the processes in the city, it has to be experienced with all senses. The Live Geography approach presented within this thesis does not consider the term "senses" a solely human uniqueness, but also suggests the integration of non-human senses for experiencing the city.

Summarising, the significance of the described definitions of the term city for this thesis is that the city as a functional construct shall not only be seen as a place of social interactions, financial transactions, a network of technology nodes, a geographical agglomeration area or a political landscape, but more as an actuated multi-dimensional conglomerate of heterogeneous processes, in which the citizens are the central component. In other words, the city can be seen as a complex near real-time control system (Resch et al., 2008), creating a feedback loop between the city itself, the city management and the citizens, which is achieved by ubiquitous sensing of a variety of parameters.

3.3 Particular Challenges of Urban Monitoring

The urban context poses many challenges to pervasive monitoring: apart from well-known general monitoring challenges, very particular issues arise for geo-sensor network deployments in the city. These range from physical sensor mounting to social and privacy implications. Furthermore, the sensitive urban political landscape has to be accounted for, which might cause unforeseen challenges.

3.3.1 Technological and Technical Issues

The key technological challenge is the integration of different data sources owned by governmental institutions, public bodies, energy providers and private sensor network operators. This problem can be tackled with self-contained and well-conceived data encapsulation standards – independent of specific applications – and enforced by legal entities, as discussed in sub-chapter 3.6. However, the adaptation of existing sensors to new standards is costly for data owners and network operators in the short term, and so increased awareness of the benefits of open standards is required.

From a technical viewpoint, unresolved research challenges for ubiquitous urban monitoring infrastructures are manifold and include: finding a uniform representation method for measurement values, optimising data routing algorithms in multi-hop networks, data fusion, and developing optimal data visualisation and presentation methods. The latter issue is an essential aspect of decision support systems, as different user groups might need different views of the underlying information. For example, in case of emergency local authorities might want a socio-economic picture of the affected areas, while first-response forces are interested in topography and people's current locations, and the public might want general information about the predicted development of a disaster.

Naturally, there are a number of well-known domain-specific technical issues (energy supply, mote size, robustness, routing, ad-hoc network connections, reliability, connectivity, self-healing mechanisms etc.). These have to be addressed as the case arises depending on specific requirements of the end application.

Moreover, highly unpredictable challenges exist arising from the openly accessible, dynamic and variable urban environment, such as severe weather conditions, malfunctioning hardware, connectivity, or even theft and vandalism. These general and mostly technical issues are well elaborated and shall not be discussed within this thesis, as elaborated in more detail in sub-chapter 2.4.

3.3.2 Various Stakeholders

Other – often neglected – issues for urban sensor network deployment are socio-political concerns and thematic challenges. This means that the feedback loop depicted in Figure 3 has to be taken into account when designing urban sensor networks. In practice, various kinds of stakeholders have to be considered including citizens, information providers, research institutions, politicians, the city management, or other influential interest groups. This cycle involves all steps of the deployment process from planning, sensor placement, customised information provision, and feedback from the citizens and other interest groups (Resch et al., 2009b).




Figure 3: Feedback Loop for Environmental Monitoring in the City.

From a more contextual standpoint, an important peculiarity of the urban context is that there are large variations within continuous physical phenomena over small spatial and temporal scales. For instance, due to topographical, physical or radiometric irregularities, pollutant concentration can differ considerably, even on opposite sides of the street. This variability tends to make individual point measurements less likely to be representative of the system as a whole. The consequence of this dilemma is an evolving argument for environmental regulations based on comprehensive monitoring data rather than mathematical modelling, and this demand is likely to grow.

Consequently, the deployment of a large number of sensors allows for more representative results together with an understanding of temporal and spatial variability. One way to overcome this issue is to “sense people” and their immediate surroundings using everyday devices such as mobile phones or digital cameras, as proposed by Goodchild (2007). These can replace – or at least complement – the extensive deployment of specialised city-wide sensor networks. The basic trade-off of this people-centric approach is between cost efficiency and real-time fidelity. The idea of using existing devices to sense the city is crucial, but it requires more research on sensing accuracy, data accessibility and privacy, location precision, and interoperability in terms of data and exchange formats. Furthermore, measurements are only available in a quasi-continuous distribution due to the high spatial and temporal variability of ad-hoc data collection. Addressing this issue will require complex distribution models and efficient resource discovery mechanisms in order to ensure adaptability to rapidly changing conditions.

3.3.3 Impact and Privacy Implications

Another central issue in deploying sensor networks in the city is the impact of fine-grained urban monitoring, as terms like “air quality” or “pollutant dispersion” are only surrogates for a much wider and more direct influence on people, such as life expectation, respiratory diseases or quality of life. This raises the demand of finding the right level of information provision. More accurate, finer-grained or more complete information might in many cases not necessarily be worthwhile having, as this could allow for drawing conclusions on a very small scale, in extreme cases even on the individual. This again could entail a dramatic impact in a very wide range of areas like health care, the insurance sector, housing markets or urban planning and management.

This raises the challenge to find the balance between providing pervasive real-time information while still preserving people’s privacy. Public concerns about unauthorised monitoring and possible use of these publicly available data streams in ways that would impinge upon personal privacy or be unfairly beneficial to a commercial interest with no compensation to the public have to be addressed carefully.

Therefore, it is self-evident that the provided information has to be highly accurate, reliable and unambiguous. Thus, quality control and error prevention mechanisms including appropriate external calibration are even more important for monitoring networks in the city than in other, less connected, environments. This may create

pressure for new enforcement actions on the part of state or federal agencies with regulatory responsibilities.

Finally, some more unpredictable challenges posed by the dynamic and volatile physical environment in the city are radical weather conditions, malfunctioning hardware, connectivity, or even theft and vandalism. Moreover, there are a number of rather obvious but non-trivial challenges to meet, such as optimal positioning of sensors, high spatial and temporal variability of measured parameters or rapid changes in the urban structure, which might cause considerable bias in the measurements.

3.4 Sensor Web Technology

Regarding the generation of measurement data, the creation of the Live Geography approach requires pervasive sensing technology in an intelligent infrastructure. Over the last few years, sensor web has been a rapidly emerging technology, which serves for gathering measured data and combining them to generate an overall combined result.

For the formation of such a sweeping vision of a “digital skin for planet earth” as envisaged by Neil Gross (1999), and described in the initial paragraphs of this introduction, sensors have to be inter-linked and easily accessible to provide live sensor data. This enables near real-time decision support in diverse areas such as environmental monitoring, medical emergencies, disaster management or public safety. The following figure illustrates the principal concept for the creation of such an Earth-spanning system.




Figure 4: Basic Sensor Web Concept. (Delin et al., 2005)

In view of technical requirements, sensor webs have to fulfil several criteria such as interoperability, intelligence and scalability, which will be discussed in more detail in sub-chapter 5.1. Although the combination of different types of sensors is a prescription originating from the definition of a sensor web, a number of proprietary high-end sensor network solutions have been developed recently. These are surely conformant to underlying IP-based communication standards, but do not possess higher-level compliance, i.e. defined data structures, service interfaces, data provision mechanisms or standardised alerting capabilities. Thus, the concept developed within this thesis implies standardisation and openness as a premise in order to guarantee broad applicability.

Generally speaking, sensor webs are a network of intra-communicating sensors, which might be of a different type, used to monitor the environment. Monitored parameters are manifold, including temperature, precipitation, different substances in the atmosphere, processes within the human body, industrial control functions, and innumerable more, as shown in Figure 5.




Figure 5: Sensor Web with Inter-communicating Sensors. (Liang et al., 2004)

An important aspect of sensor webs are interoperability of sensors, data structures and measurement provision, as elucidated in sub-chapter 3.4. Thus, the Open Geospatial Consortium (OGC) started the Sensor Web Enablement (SWE) initiative in 2003, which aims to make sensors discoverable, accessible and controllable over the internet in a standardised manner. SWE is described in more detail in sub-chapter 6.2.

3.5 Why Standards?

Taking the challenges of urban monitoring networks into account it is apparent that standardisation is a vital prerequisite for the creation of interoperable and portable infrastructures. To achieve maximum interoperability, standardisation has to be applied throughout the whole workflow, ranging from unified data encapsulation and representation to service interfaces for data provision, event detection and alert generation, and even to workflow management.

3.5.1 Standardisation Initiatives

As shown in Figure 6, there are a variety of standardisation efforts by different institutions aiming at interoperability of (geospatial) data and services. Basic infrastructure-related standards – such as web service or communication standards as specified by the World Wide Web Consortium (W3C) – are mostly “de facto” standards (commonly used technological specifications). On the contrary, domain-specific definitions are mostly legally binding (“de jure”). As de facto standards are often developed “on-demand”, i.e. by a specific group of mostly industrial companies, and norms are often imposed by legislation, the Open Geospatial Consortium (OGC) tried to bridge these two worlds. This happens by creating service interfaces using established de facto standards while still being open for incorporating legally binding domain-specific norms.




Figure 6: Standardisation Bodies Classified by Legal Bindingness and Genericness.

The OGC – the leading institution for the establishment of geospatial standards – pursues a double-sided approach. General standards are developed in a top-down

fashion. For instance, the Geographic Markup Language (GML) standard has been established in a top-down manner trying to model every single object in the world while being as generic and comprehensive as possible.

In contrast, standards such as the Keyhole Markup Language (KML) have been created out of a concrete need of specialised companies (in the case of KML, Google needed a format for data representation to be displayed in their 3D viewer Google Earth). In case of broad support and development perspectives, these formats are then officially adopted as OGC standards.

3.5.2 Benefits Arising from Usage of Standards

The establishment of an open system architecture can benefit in various ways from the use of open standards. First, standardised methods of representing data, providing data, and visualising information foster interoperability between services and different heterogeneous data sources. This is an essential prerequisite for deploying open systems like the Live Geography infrastructure for the use of various types of interest and user groups. Furthermore, different end users can be served by very intuitive and tailored graphical interfaces by decoupling data and visualisation.

Another crucial benefit of using open standards is simple extensibility in terms of integrating new data sources and analysis components if the system is established in a modular structure. This means that business or data analysis processes can benefit by this straight-forward integrability of additional kinds of data and data resources.

The underlying enabling process of the above mentioned benefits is that standards allow for automated machine-to-machine (M2M) communication. Computers can communicate via well-known interfaces and thus combine various services. This is also referred to as service-chaining in Service Oriented Architecture (SOA) jargons. This M2M-enabled communication infrastructure results in wide reusability of different software components, which can be easily adapted to perform different tasks.

Furthermore, uniform interfaces guarantee request consistency, which facilitates the implementation of client applications. Subsequently, well-established geospatial interfaces such as OGC Web Feature Service (WFS), Web Map Service (WMS) or Web Coverage Service (WCS) ensure mid-term preservation of request structures in contrast to custom Application Programming Interfaces (API), which can change frequently due to newly arising specific application needs. Those interface standards are elaborated in more detail in sub-chapter 6.2.

One trade-off, which has to be accounted for when using open standards, is that increased effort is needed to establish the system when first deploying it. The time-consuming part is to implement standards, which can be widely avoided in the GIS domain by using open-source software where applicable.

However, on the long-term, operation, administration and transaction costs are minimised using open standards as maintenance and communication are much simpler and cheaper, as stated in DIN German Institute for Standardization (2000) and Anderson (2004).

Summarising, it can be stated that open standards in urban measurement infrastructures foster interoperability and reusability. Especially the use of widely established OGC interfaces such as WFS, WMS or WCS is a strong basis for an extensible system infrastructure. The main goal is to achieve the highest possible degree of consistency within the whole system in terms of request uniformity.

Moreover, the integration of new and emerging standards such as OGC Sensor Web Enablement (SWE) is a main focus in the data gathering part of the Live Geography approach. This allows for automated M2M communication as sensor data can be automatically integrated into real-time analysis processes.

3.6 Policy-frameworks for the Integration of Real-time Sensor Information

During the past several years, we have witnessed an explosion of spatial data collection and availability in digital form. There are various national and international efforts to establish spatial data infrastructures (SDI) for promoting and sharing geospatial information throughout governments, public and private organisations and the academic community.

It is a substantial challenge to solve political, technological and semantic issues for sharing geographic information to support decision making. In 2007, the United Nations Geographic Information Working Group published a report subsuming recent regional, national and international technologies, policies, criteria, standards and people necessary to organise and share geographic information. The basic aim is to include real-time location aware sensor measurements to develop a United Nations Spatial Data Infrastructure (UNSDI) and encourage interoperability across jurisdictions and between UN member states.

These SDIs are intended to stimulate the sharing and re-use of expensive geographic information by several initiatives:

- On a supra-national level, the INfrastructure for SPatial Information in Europe (INSPIRE) aims to enable the discovery and usage of data for analysing and solving environmental problems by overcoming key barriers such as inconsistency in data collection, lack of documentation, and incompatibility between policies and existing geographic information systems (GIS). (Commission of the European Communities, 2009a)
- Global Earth Observation System of Systems (GEOSS) seeks to establish an overarching system on top of national and supra-national infrastructures to provide comprehensive and coordinated earth observation. The goal is to transform these data into vital information for society. (US Environmental Protection Agency, 2009)
- Global Monitoring for Environment and Security (GMES) can be considered a European instantiation of the GEOSS programme. GMES is a European initiative for the implementation of information services dealing with environmental and security issues using earth observation information and in-situ data for the short, mid and long-term monitoring of environmental changes. This is Europe's main contribution to the Group of Earth Observations (GEO) for monitoring and management of planet earth. (Commission of the European Communities, 2009b)
- The Shared Environmental Information System (SEIS) is an overarching approach by the European Union to modernise and simplify the collection, exchange and use of environmental data and information to allow for more accurate and informed environmental policy-making. INSPIRE and GMES – amongst others – are concrete initiatives towards the realisation of SEIS. (Commission of the European Communities, 2008)
- The Global Spatial Data Infrastructure Association is one of the first organisations to promote international cooperation in developing and establishing local, national and international SDIs through interaction between organisations and technologies supported by the U.S. Geological Survey (USGS). (Onsrud, 2009)

The common goal of all these initiatives – and the numerous national SDI approaches – is to integrate and share geospatial information comprising real-time measurement data sets, remote sensing data and environmental information. With the European Shared Environmental Information System (SEIS), a new concept has been introduced to collect, analyse and distribute these information sets in a loosely coupled, “federal” structure. A particular focus is dedicated to the definition of interfaces for the provision of near real-time datasets like sensor measurements.

This new flexibility should foster the integration of sensor measurements in existing SDIs.

3.7 Use Cases and Application Scenarios

As mentioned in chapter 2, the main aim of this thesis is the development of a standardised technological infrastructure, upon which several possible applications can be implemented. In order to be able to create a new innovative vision for this generalised architecture, several implemented use cases (with known shortcomings) and visionary application scenarios are examined up-front.

3.7.1 Use Cases

A number of urban sensor applications already exist, which are mostly built up in monolithic systems. The introduction of the Live Geography infrastructure can possibly yield a significant performance increase and create a high level of interoperability, most possibly widening the impact field of these applications.

I Am Alive System

Tada et al. (2000) describe the "I Am Alive" (IAA) system, which collects data about people's physical condition via a web interface, for instance after an environmental disaster. Apart from collecting people's personal information, the system also offers vital information such as water and food availability, evacuation and medical care sites and traffic conditions.

Originally, the IAA system arose from the Widely Integrated Distributed Environment (WIDE) project, whose goal was to build a robust distributed network. To achieve this vision, the "Lifeline Working Group" was organised to investigate the Internet's availability and reliability in case of a disaster. More information about the project can be found in WIDE Project (2005).

The concept of IAA was created taking into account the key criteria scalability and flexibility, robustness and fault tolerance, privacy and security, as well as data accuracy. In the course of the project, the project consortium conducted five disaster drills, which aimed at preparing people for Internet use in case of a catastrophe. The goal was to gain user experiences about their requirements concerning Internet use in case of a disaster and entering the results into the IAA backend system. During the five drills, different technologies such as mobile phones, fax, telephone, the Internet and e-mail were used as input interfaces.

The results of the drills, during which up to 10.000 queries were posted per trial, yielded that additionally to the original design issues (daily operation in the sense of reliability and faultlessness, internationalisation by the support of different languages and inter-connectivity with other systems, meaning a standardisation of the protocol for exchanging data among them), several other critical points have to be treated. Firstly that is to face mobility issues, i.e. to create a plug-and-play function; secondly to build a more general framework in order to support applications such as multi-cast; and finally to fulfil security and privacy requirements in such a discerning environment.

After numerous public trials, interconnectivity tests, and standardisation efforts, the IAA system has been used officially for the first time after the earthquake, which occurred on 20 March 2005 in Fukuoka, Japan. More information about the IAA system can be found at the alliance's homepage (IAA Alliance, 2005) and in Tada et al. (2000).

Water Recharge Basin Monitoring

In November 2003, 16 sensing devices were deployed by the National Aeronautics and Space Administration (NASA) at a water recharge basin in Arizona to observe flooding events and compare the results to satellite-based measurements. (Delin et al., 2005)

The sensor devices measured the movement of water across the basin's surface and observed the penetration of water into the soil. The pods contained sensors for measuring air and soil temperature and humidity as well as the light level. The pods were organised in a distributed network, where the sensor nodes communicated wirelessly with each other.

The authors state that the most critical points concerning sensor webs are the pods' power consumption in connection with battery size, the antenna size, transducers and the ease of deployment, which is tightly coupled with sensor size, weight, robustness, and needed effort for establishing the network.

More detailed information about this sensor web project, which was finally decommissioned in April 2005, can be found in Delin et al. (2005) and Delin and Woodrow (2005).

3.7.2 Application Scenarios

There are a wide variety of possible application scenarios for the Live Geography approach, which have not been feasible until now due to technical and conceptual

shortcomings. The following sub-chapters describe a broad spectrum of future urban monitoring applications.

The Smart City

The Smart City project, which is a joint long-term initiative between IBM and the City of Salerno, Italy, aims to identify and design an innovative model for a sustainable city. This shall happen through a range of digital technologies, adding intelligence to existing infrastructure. The goal is to optimise resources, improve services to citizens and businesses, and make the city “smarter” and more efficient. (Questex Media Group, Inc., 2009)

The project builds upon four key initiatives; first, to place smart sensors throughout the city to overcome architectural barriers and to support visually impaired people in their navigation through the urban environment; second, to deploy smart tags at places of cultural and artistic heritage, and thus supply tourists with location-relevant information through their mobile phones on-the-go; third, to provide elderly people with a simple electronic device for “remote” health monitoring; and fourth, to use existing technological infrastructure data, such as traffic light status, parking space availability or traffic counts, to manage and share traffic information in real-time.

It is evident that a key requirement for this development of a Smart City is interoperability between sensors, monitoring systems, gathered data and end user devices. The Live Geography approach tries to tackle these issues by providing a fully standardised measurement infrastructure.

Public Health Monitoring

Public Health has been asked to participate in policymaking on “quality of life” issues increasingly over the past decade. The superimposing of the medical model to describe the impact of conditions that have traditionally been regarded as nuisances has created a great challenge, particularly in the field of environmental health. (Resch et al., 2009b)

One pollutant often used to serve as a proxy is NO_x, which technically represents various gaseous species comprised of oxygen and nitrogen molecules. Another indicator of near-roadway effects that has gained recent attention is ultrafine particulates (UFP), particles that are less than 0.1 microns (100nm) in diameter. Thus, air quality measurements of hazardous air pollutants can be widely associated with traffic (non-point sources). A pervasive sensor network could help capture measurements in high spatial and temporal resolution to take short-term

measures (dynamically adapt traffic management or send out alerts to citizens in case of threshold exceedance). Also, it could support traditional long-term studies on the impact of certain pollutants on public health.

A practical use case scenario for the developed Live Geography approach in the area of air quality monitoring could look like as follows. In the first step, various CO and ozone sensor measurements are fused by a Sensor Observation Service (SOS), i.e. data from several sensor networks are harmonised into the common Observations and Measurements (O&M) data format in a sensor fusion process. Next, different SOS services – possibly operated by different authorities – are combined in a data fusion process. This implies the harmonisation of the different measurement data by applying methods for fusing heterogeneous data, like calibrating earth observation with in-situ CO₂ sensor fusion results, time representation adaptation, and unit conversion. Then, these data are fed into a geo-processing operation, which combines them with other geospatial layers: census data, a road network and recreational areas). In this way, a city-wide fine-grained air quality indicator index is created. The applied kriging algorithm is likely to reveal previously unseen correlations between these parameters, and is therefore the basis for spatial decision support for the health department of the city.

Noise Mapping

The use case of noise mapping in urban environments has received a lot of attention recently. Many disputes within the research field emerge from noise impacts associated with construction, excavation or some other commercial or industrial enterprise. (Resch et al, 2009b)

These disputes also originate from use of domestic landscaping equipment, like leaf blowers and snow blowers. The limits imposed by the city on noise generation are intended to assess the background noise levels. A source cannot be held responsible for noise levels that exceed the city's allowable limits if the ambient noise in that area already exceeds those limits.

The development of noise “maps” may not immediately result in satisfaction from aggrieved residents, but it can be used to consider the noise impact of future development and zoning policies. It may also contribute to efforts to reduce the number of cars travelling across the city by adding the noise impact dimension to the discussion. This is much more likely to be given full consideration if it can be demonstrated with highly resolved data maps, which can be generated in near real-time using the Live Geography infrastructure.

Urban Heat Island Detection

The urban heat island effect describes the contribution of the built environment to the ambient temperature within urban areas. While this is not likely to become a primary public health concern, it has great bearing on efforts to limit the loss of heating energy across the city.

Different agencies have been established to work on a long-term strategy to reduce overall energy use (e.g. Cambridge Energy Alliance¹) and to encourage individual homeowners and building owners to evaluate their energy loss. It is quite possible that small changes in heat loss, as described through a detailed heat map of the city over time, could show progress towards energy efficiency in a materials way. This could be used both as an evaluation tool in tracking the city's progress, and as a means to engage the public in the energy goals of the community.

Smart Dust Applications

Steel (2005) describes "smart dust", a new technology aiming at reducing the size of wireless sensor motes as far as possible. Smart dust pods normally comprise environmental sensors to measure parameters such as temperature, light, humidity or position. Furthermore, also additional technologies such as Radio Frequency Identification (RFID) tags can be embedded.

Depending on the size of the pods, capabilities and longevity can be increased. A highly miniaturised sensor pods is depicted in Figure 7. The "Spec" mote comprising full sensing and communicating technology without its battery is about the size of an aspirin tablet.




Figure 7: "Spec" Sensor Pod next to the Tip of a Ball Point Pen. (Steel, 2005)

¹ <http://www.cambridgeenergyalliance.org>

An operating system, which was especially designed for these small wireless embedded sensor motes at UC Berkeley, is called TinyOS. This open-source operating system, which is running on smart dust sensor pods, is described in UC Berkeley (2009).

Smart dust sensor pods can be arranged in a self-organising manner to retain the network's extensibility. The gathered data are collected by a base station and transmitted to a central server via a cellular link and are employed in a specific workflow to dynamically generate reports, alerts and developing graphs, which are automatically updated via web services. More technical information about smart dust can be found in Warneke et al. (2001).

The advantages of such a system are the possibilities of controlling urban processes in real-time (e.g. traffic management or energy load balancing), and to combine measured data with other data sources, e.g. to correlate environmental data with weather forecasts and traffic emergence.

The main challenges of building such smart dust implementations are hardware protection against physical damage, accomplishment of a reliable communication link, and to build the adequate application logic to generate useful output. The Live Geography approach addresses these issues by providing a modular architecture interlinking components by open standards.

International Emergency Alert System

Baker and Carpenter (2005) propose a conceptual approach for an Emergency Alert System (EAS), which originated after a devastating earthquake near Sumatra on 26 December 2004. This proposal builds on currently existing technologies such as the Internet, e-mail, Short Message Service (SMS) broadcast etc. The special innovation of the EAS is that alerting messages are targeted solely to the affected area.

The implementation proposal of Baker and Carpenter requires four main system components. Firstly, this is a warning centre, which monitors appropriate parameters such as seismic activity. The installation of such a warning centre consists of a specific selection of possibly affected regions and the cross-linking with the needed sensor network. Secondly, a mailing list of people, who are interested in alerts and connected to the Internet, is required in order to convey warning messages via authenticated e-mail or the World Wide Web (WWW), which is realised via sufficiently frequent polling, e.g. by using Really Simple Syndication

(RSS)². Thirdly, a clear warning interpretation policy has to be established so that receivers know how to interpret the warning messages and how to react. Lastly, the created alerts should finally be distributed to all affected people as in a first step, only a limited number of governmental or private-sector service institutions receive the messages.

Generally speaking, the main objective concerning the message distribution is the conveyance of precise, understandable and timely alert messages to people who might be affected. Potential technologies to broadcast the warning messages include mobile telephone text message services, broadcast services like television (TV) or radio, presence-based services and message-based services.

Mobile telephone text message services comprise the Global System for Mobile Communications (GSM) SMS and similar services provided by other systems such as Code Division Multiple Access (CDMA) or other proprietary technologies. As SMS, whose messages can carry textual content of up to 160 characters, is available in wide parts of the world and as the whole service is location-aware, it forms a solid data transmission basis for an international alerting service.

Broadcast services such as television or radio could be interrupted in case of an emergency to broadcast an alert message very quickly. However, not all affected people pay whole attention to the broadcasted programme and might thus miss the alert.

Both presence-based and message-based services require a subscription mechanism. For presence-based services, this means that a person has to subscribe an alert service for a specific region. If an emergency occurs and the user is currently located at an endangered region, the person receives a warning message by the system. For message-based services, the subscription system is principally the same except that the user receives an authenticated e-mail message for a certain region regardless of his or her current location.

Naturally, during the conception of such an alert system, security considerations have to be made because if the system is intruded, alerts might be spoofed decreasing the reliability of the whole service. Thus, in order to retain accuracy and the degree of reliance, alert messages should be authenticated and the mailing lists should be periodically tested to ensure timeliness.

² RSS is an XML-based information system containing short descriptions of web content to keep automatically up-to-date with recent news.

Future Earth Observing System

A very ambitious, but nonetheless realistic long-term vision of a global earth observing system is expounded by Hildebrand et al. (2003). The authors describe both the requirements and the architecture of an internationally deployed sensing system. This system is not particularly targeted to urban environments, but shows the broad vision of an interoperable and ubiquitous sensor network.

The long-term goal is a sensor web-based observation architecture comprising innumerable remote and in-situ sensing devices operating directly at the Earth's surface or from Geostationary Orbit (GEO) or Low Earth Orbit (LEO) satellites, whereby the former could serve for collecting measurements with high temporal update rates up to hourly and the latter could be employed for sensing processes occurring only in longer intervals.

Parameters, which should be frequently measured, are aerosols, chemical substances in the atmosphere as well as precipitation and wind characteristics. Examples of rather seldom gauged factors are ocean evaporation, ice processes, crustal deformation as well as the physiological state and dissolved organic matter.

The final vision of this approach is to explore connections between different processes on Earth such as the linkage between the climate and the water cycle in the oceans.

With these measured parameters, several processes on Earth could be comprehended and predicted, respectively in a next step. First, the combination of measurements of temperature, precipitation, sea ice thickness and water availability is necessary to learn about the intra-seasonal climate. This would allow for monthly or even yearly climate prediction. However, this forecast would require the maintenance of an accurate history and the precise calibration of the instruments over a very long time period.

Secondly, extreme weather conditions such as tropical storms could be predicted by combining sea surface temperature, water vapour and cloud-drift winds with adequate statistical models involving a high resolution and a high update rate, which could be realised by a network of LEO spacecrafts.

The next challenge is the prediction of sea level changes taking into account causative factors such as ice sheet thickness and melting, water cycle fluctuations and the oceans' volume variability. However, to be able to draw conclusions concerning these processes, numerous daily, monthly and annual observations such as soil moisture, ice sheet topography and motion or ocean bathymetries have to

be made. This necessitates an accurately calibrated system over a long period and a measurement precision similar to current in-situ sensing.

Another application of a global earth observing system is the investigation of the linkage between biosphere-climate interactions and human influences due to urbanisation. This requires knowledge about biosphere-land-ocean exchange, natural regulatory controls and effects of human behaviour. These complex processes can be derived from the combination of several single measurements such as atmospheric aerosols, radiation, pollution, precipitation or assessments of the standing biomass.

Finally, predictions on earthquakes and water availability could be made with a global earth observing system. The former requires accurate pre-, co-, and post-seismic observations in all phases of the earthquake, whereas the latter accounts e.g. for global precipitation, vapour transport, soil moisture or ocean circulations.

Conclusion

As it can be seen from the above fictive future use case descriptions, the realisation of a city-wide or even earth-spanning observation system will be a complex process that has to be carried out successively in a world-wide common initiative, as the system architecture is composed of several different satellite systems and measurement instruments. However, system designers also have to take into account the parallel development of measurements and observations, scientific comprehension and computer models. Furthermore, the long-term calibration has to be retained considering evolving measurements and predictive models.

Chapter Summary

This thesis presents the "Live Geography" approach, which proposes a standardised infrastructure for monitoring applications with a particular focus on urban environments. Before addressing technical questions, the establishment of such an architecture requires a definition of the term "city" as a functional construct. For the approach developed within this thesis, the city shall not only be seen as a place of social interactions, financial transactions, a network of technology nodes, a geographical agglomeration area or as a political landscape, but more as an actuated multi-dimensional conglomerate of heterogeneous processes, in which the citizens are the central component. In other words, the city can be regarded as a complex near real-time control system, creating a feedback loop between the city itself, the city management and the citizens, which is achieved by ubiquitous sensing.

This urban context poses many challenges to pervasive real-time monitoring: apart from well-known general monitoring challenges, very particular issues arise for geo-sensor network deployments in the city. These range from physical sensor mounting to social and privacy implications. Furthermore, the sensitive urban political landscape including heterogeneous interests of a variety of stakeholders has to be accounted for, which might cause unforeseen challenges.

It shall be noted that the strict term "real-time" in the context of this thesis can be interpreted as "at present" for urban monitoring applications. This means that the aim is to assess the environment "now", not a historical and perhaps outdated representation. To account for this non-rigorous requirement, the term "Live Geography" seems better suited than "Real-time Geography". In this reflection, "near real-time" appears to be closest to „live", as it does not impose rigid deadlines and the expression itself suggests dynamic adaptation of a time period according to different usage contexts.

However, the ambitious goal of establishing such a highly flexible, interoperable, modular and portable measurement infrastructure requires the extensive use of open geospatial standards (for geo-sensor networks and data provision) and according policy frameworks such as INSPIRE, GMES, SEIS, GEOSS etc. are vital prerequisites.

RELATED WORK

RELATED WORK

4 Related Work

The Live Geography approach is manifold in terms of both concepts and employed technologies. As such, there are several research initiatives that form part of the overall methodology. This chapter comprises a review of recent approaches towards the establishment of measurement infrastructures. The respective sub-topics of *urban sensor networks*, *real-time data integration* and *web-based geo-processing* are discussed in a separate chapter 5.

4.1 Sensor Networks and GIS

There are a number of approaches to leveraging sensor information in GIS applications. Kansal et al. (2007) present the SenseWeb project, which aims to establish a Wikipedia-like sensor platform. The project seeks to allow users to include their own sensors in the system and thus leverage the “community effect”, building a dense network of sensors by aggregating existing and newly deployed sensors within the SenseWeb application. Although the authors discuss data transformation issues, data fusion, and simple geospatial analysis, the system architecture is not based on open (geospatial) standards, only standard web services. The web portal implementation, called SensorMap, uses the Sensor Description Markup Language (SDML), an application-specific dialect of the OGC SensorML standard.

Paulsen (2008) presents a sensing infrastructure that attempts to combine sensor systems and GIS-based visualisation technologies. The sensing devices, which measure rock temperature at ten minute intervals, have been developed focusing on optimising resource usage, including data aggregation, power consumption, and communication within the sensor network. In its current implementation, the infrastructure does not account for geospatial standards in sensor observations. The visualisation component uses a number of open standards (OGC WMS and WFS) and open-source services (UMN Map Server, Mapbender).

Another sensing infrastructure is described by Murty et al. (2008). The CitySense project uses an urban sensor network to measure environmental parameters and is thus the data source for further data analysis. The project focuses on the development of a city-wide sensing system using an optimised network infrastructure. An important parallel with the work presented in this paper is that CitySense also considers the requirements of sensor network setup in an urban environment.

A GIS mashup for environmental data visualisation is presented in the nowCOAST application (National Oceanic and Atmospheric Administration, 2008). Data from several public providers are integrated in a web-based graphical user interface. nowCOAST visualises several types of raw environmental parameters and also offers a 24-hour sea surface temperature interpolation plot.

The most striking shortcoming of the approaches described above and other related efforts is that their system architectures are at best partly based on open (geospatial) standards.

4.2 Open Sensor Data Integration Architectures

The next related research area is the development of an open data integration architecture in a non-application-specific infrastructure. Recent research efforts focus on general concepts in systems architecture development and data integration, but there are mostly no concrete conclusions as to how to establish such an infrastructure. A more technical approach for ad-hoc sensor networks is described by Riva and Borcea (2007), where the authors discuss application-motivated challenges of combining heterogeneous sensor measurements through flexible middleware components. The method is strongly application-centred and thus very well-thought-out as far as specific implementation details are concerned.

Sensors Anywhere (SANY), which is presented by Douglas et al. (2008) aims to contribute to the joint European Commission (EC) and European Space Agency (ESA) "Global Monitoring for Environment and Security" (GMES) initiative by improving the interoperability of in-situ sensors and sensor networks. Its aim is also to allow quick and cost-efficient reuse of data and services from currently incompatible sources in future environmental risk management applications. This ambitious goal shall be achieved by (1) specifying a (generic) standard open architecture for fixed and moving sensors and sensor networks capable of seamless "plug and measure" and sharing (virtual networks); (2) developing and validating re-usable data fusion and decision support service building blocks and a reference implementation of the architecture; (3) closely working with end users and international organisation in order to assure that the outcome of SANY contributes to future standard(s) applicable to GMES. In contrast to the Live Geography approach presented within this thesis, SANY tries to create a very general architecture whereas this thesis focuses on urban environments. Furthermore, Live Geography – in addition to sensor network applications – also investigates the coupling with sensor fusion, alerting and server-based processing mechanisms and technologies.

4.3 Empirical Aspects – A State of the Art Review

From the establishment of former sensor web architectures, several empirical aspects can be retrieved, which should be considered in the conception of a generalised technical basis and the implementation of a generic solution.

4.3.1 Limitations

The first technical confinement of geo-sensor networks is the maximum amount of transmitted data and the processing steps in a sensor pod, which have to be investigated in a continuative investigation taking into account the requirements of the actual application. This involves publishing sensor data via a web service, which can directly access the sensors' measurements. Furthermore, the expenses of the sensor pods play a key role considering that the costs for the pods rise proportionally to its intelligence and processing capabilities. Moreover, the pod's energy consumption is a question to clarify as it rises with increasing intelligence.

Naturally, strict real-time requirements are always a limiting factor, whereby the demands in the implementation in question are varying depending on the application context, as explained in sub-chapter 3.1. Apart from strict definitions, the term "real-time" in the context of this thesis can also be expressed as an "exact point in time", which is the same for all sensors to create a significant measurement outcome. Furthermore, it can also be regarded as the possibility to start a synchronous conversation at a certain time, which might e.g. be important for geographical applications to enable the generation of a parameter development graph over a defined period of time.

4.3.2 Technical Deliberations

When designing a sensor network, several technological choices have to be preconceived, e.g. the decision for a heterogeneous or a homogeneous network, single-hop or multi-hop communication or the Medium Access Control (MAC) mechanism. These deliberations are only superficially discussed at this point as they are not of major importance for the achievement of the goals of this thesis.

Homogeneous vs. Heterogeneous Networks

In Mhatre and Rosenberg (2004), a comparative study between two sensor networks with homogeneous and heterogeneous sensing devices, respectively, is presented. In this case, homogeneous means equality in battery energy and hardware complexity. In homogeneous networks, periodic rotation is used so that each of the identical sensor pods acts as a cluster node for approximately the same

time, whereas in heterogeneous networks, there are two types of nodes, namely pure sensor nodes and cluster heads.

Mhatre and Rosenberg state that in general, homogeneous networks should be used in order to guarantee a lifetime, which is approximately the same for all pods. This fact ensures that not too much unused energy is wasted when the sensor web is abandoned. However, heterogeneous networks can be more cost effective when deploying a large amount of pods because not all sensor devices need complex hardware capable for long range transmissions with the base station. Nevertheless, the authors constitute that homogeneous networks are more robust as a single node failure does not paralyse the whole network.

Single-hop vs. Multi-hop Communication

Furthermore, the advantages of a single-hop and a multi-hop network have to be examined. The former technology is a network, in which the sensor nodes communicate directly with the cluster head, which then sends the aggregated data to a base station. In a multi-hop network, the pods pass information from node to node until it reaches the cluster pod. Generally, it can be stated that the multi-hop technique might be the better choice for urban sensor networks, in which there is a big distance between the pods, because energy consumption increases exponentially with the distance. (Mhatre and Rosenberg, 2004)

4.3.3 External Dependencies

When planning an application, which superimposes on a basic concept to form a meaningful prototype, a primary aim is to avoid external dependencies as far as possible. This means that the project progress should be independent of the good will of other organisations. If external dependencies exist, the implementation has to respond to dedicated needs and should be documented by a clearly formulated contract to guarantee orderly accomplishment of the project.

A second avoidance of external dependencies originates from the mentioned fact that sensors can be queried separately and directly via a web service application. This eliminates the dependency on a central data server, which can often form a bottle neck when many external applications have to be served with vast amounts of data.

4.3.4 Security Aspects

The examination of security issues implies encryption of the transmitted data, authentication of users to retrieve data and to guarantee a certain degree of

reliability of the system in terms of correctness and availability. As mentioned in sub-chapter 2.4, security requirements vary depending on the end application. Thus, the development of a generic security framework is not a vital part of this thesis.

Data can be easily encrypted when using web services based on the Simple Object Access Protocol (SOAP) using the Web Service Security (WSS) standard, which offers an encryption mechanism by itself making use of security tokens in connection with digital signatures as described in sub-chapter 4.3. Moreover, intra-sensor communication can be encoded, which is achieved by availing the particular encryption standard provided by the used transmission technology. User authentication can for instance be realised by a Universal Unique Identifier (UUID) for every user. Finally, the system's reliability can be increased by consistent programming and careful planning of the implementation.

However, before being able to create a security solution for a wireless sensor network, different types of attacks and precautions have to be examined.

Attacks and Precautions

As communication in a wireless sensor network cannot be physically protected, security mechanisms are of significant importance. Principally, attacks on such networks can be divided into outsider attacks, in which the attacker is not an authorised participant of the network, and insider attacks. The former one can be performed by eavesdropping the network's radio frequency, by inserting altered or spoofed packets into the network or by disabling sensor nodes, which can happen by physically destroying them or by Denial of Service (DoS) attacks, which means sending masses of useless packets to a node to bring down its communication link or to drain its battery. Insider attacks can be conducted by using compromised nodes to steal confidential information or to paralyse the whole network.

These attacks could be precluded by letting the base station act as a point of trust by physically protecting it against attacks, whereby this procedure raises severe scalability concerns. Moreover, networks can be made more robust against outsider attacks by increasing redundancy, meaning to deploy more sensor pods as absolutely necessary for the measurements, or by employing standard security techniques such as cryptographic algorithms. Furthermore, access control policies could be employed at the base station to increase secrecy and thus to prevent information flows to adversaries. (Shi and Perrig, 2004)

When considering sensor network security, data encryption is a separate topic of interest. As asymmetric encryption mechanisms can be considered too expensive

for sensor networks in terms of computational requirements, efficient symmetric techniques can be employed. Currently, several effective encryption methods are developed, each of which offering different advantages. Despite encryption is a key requirement for assuring data integrity and security, it is only of peripheral interest within this thesis, as mentioned in sub-chapter 2.4.

Web Services Security

In a cooperative effort, International Business Machines (IBM) Corporation, Microsoft Corporation and Verisign, Inc. created the Web Services Security (WSS) concept, which is described by Kaler (2002). The primary goal of WSS is to enable secure message exchanges over SOAP providing a flexible set of mechanisms rather than defining fixed security protocols.

In general, two threats have to be accounted for when designing a message-based security mechanism. Firstly, a message could be read or altered by attackers and secondly, antagonists could exploit services with lacking security policies in order to guarantee forwarding of their messages. Thus, WSS employs digital signatures in combination with so-called security tokens, which is the part of the message containing the actual security mechanism and is embedded into the SOAP header.

In WSS, message protection is principally realised by two mechanisms, namely message integrity and message confidentiality. The former is accomplished by supporting Extensible Markup Language (XML) signatures combined with security tokens to guarantee that a message is transmitted without modifications, whereby also multiple signatures of several users can be employed. The latter is achieved by XML encryption in connection with security tokens, whereat in this context, confidentiality means that an antagonist is not able to read the message. Again, WSS supports additional encryption techniques and operations by several users.

The following lines show an example of how a SOAP message can be encoded by embedding security elements:

```

01  <?xml version="1.0" encoding="utf-8"?>
02  <S:Envelope xmlns:S="http://www.w3.org/2001/12/soap-envelope"
    xmlns:ds="http://www.w3.org/2000/09/xmldsig#">
03      <S:Header>
04          <m:path xmlns:m="http://schemas.xmlsoap.org/rp/">
05              <m:action>http://fabrikam123.com/getQuote</m:action>
06              <m:to>http://fabrikam123.com/stocks</m:to>
07              <m:id>uuid:84b9f5d0-33fb-4a81-b02b-5b760641c1d6</m:id>
08          </m:path>
09          <wsse:Security
    xmlns:wsse="http://schemas.xmlsoap.org/ws/2002/04/secext">
10              <wsse:UsernameToken Id="MyID">
11                  <wsse:Username>Zoe</wsse:Username>

```

```

12         </wsse:UsernameToken>
13         <ds:Signature>
14             <ds:SignedInfo>
15                 <ds:CanonicalizationMethod
16                     Algorithm="http://www.w3.org/2001/10/xml-exc-c14n#" />
17                 <ds:SignatureMethod
18                     Algorithm="http://www.w3.org/2000/09/xmldsig#hmac-
19                     sha1" />
20                 <ds:Reference URI="#MsgBody">
21                     <ds:DigestMethod
22                         Algorithm="http://www.w3.org/2000/09/xmldsig#sha1"
23                         />
24                     <ds:DigestValue>LyLsF0Pi4wPU...</ds:DigestValue>
25                 </ds:Reference>
26             </ds:SignedInfo>
27             <ds:SignatureValue>DJbchm5gK...</ds:SignatureValue>
28             <ds:KeyInfo>
29                 <wsse:SecurityTokenReference>
30                     <wsse:Reference URI="#MyID" />
31                 </wsse:SecurityTokenReference>
32             </ds:KeyInfo>
33         </ds:Signature>
34     </wsse:Security>
35 </S:Header>
36 <S:Body Id="MsgBody">
37     <tru:StockSymbol
38         xmlns:tru="http://fabrikam123.com/payloads">QQQ</tru:StockSym
39         bol>
40 </S:Body>
41 </S:Envelope>

```

As it can be seen from the lines of code above, the security part is simply integrated in the ordinary SOAP message, which begins with the XML declaration, the SOAP envelope and SOAP header tags in lines 01 to 03. The following five lines specify how the message shall be routed. Then, in line 09, the actual security tag, which contains security information for the receiver, is opened. The following three lines define the security token, which is done by a *UsernameToken* in this case, implying that the service knows the password, meaning that it is a shared secret.

Then, in lines 13 to 28, a digital signature is defined, whereby in this case, the signature is generated from the user's password, which is not a recommendable procedure in practice. Lines 17 to 20 then specify the elements of the SOAP message to be signed, namely the body element. For a more secure transmission, also the routing header and other message parts should be signed. Line 22 specifies the signature value of the data to be signed.

Lines 23 to 27 contain a location, where the security token associated with this signature can be found. After the header element has been closed, the actual body with the payload follows and finally, the SOAP envelope element is closed again.

Chapter Summary

The Live Geography approach is manifold in terms of both concepts and employed technologies. As such, there are several research initiatives that form part of the overall methodology. This chapter comprises a review of recent approaches for measurement infrastructures.

The most striking shortcoming of current geo-sensor infrastructures is that they mostly do not leverage open standards. This severely limits interoperability and flexibility of data, services and software components.

Furthermore, those approaches widely focus particularly on the development on a sensor network application for a specific use case. Thus, the coupling of sensor webs with real-time data integration, ad-hoc web-based data processing and broader workflows are decoupled research topics.

Moreover, current developments do often not account for their integration into existing systems and business logics. This degrades them to stand-alone solutions, which cannot be used in a broader context by inter-chaining them with legacy workflows.

METHODS

METHODS

5 Methods

The establishment of monitoring systems is not a trivial task. In particular, the deployment of geo-sensor networks in urban environments poses a variety of highly heterogeneous challenges. Apart from well-known general monitoring issues, very particular questions arise, which range from physical sensor mounting to socio-political and privacy implications, as shown in sub-chapter 3.3.

This chapter contains a background assessment of the essential parts of monitoring infrastructures, i.e. urban sensor networks, real-time data integration into GIS (Geographic Information Systems) and web-based data analysis technologies. It shall be noted that the latter research area is per se not a central part of this thesis, but it is necessary being treated when investigating the overall monitoring infrastructure.

5.1 Urban Geo-Sensor Networks

In the last decade, strongly decreasing sensor costs triggered intense research in sensor technologies and vice versa. This fact, together with miniaturisation efforts and rising awareness of the benefits of automated real-time sensor applications, resulted in the deployment of a number of urban sensor systems. Additionally, the use of live measurements in geospatial data analysis systems rapidly gained importance.

5.1.1 Literature Review

The first research domain treated at this point is **sensor network development for urban monitoring**. The *Oklahoma City Micronet* (University of Oklahoma, 2009) is a network of 40 automated environmental monitoring stations across the Oklahoma City metropolitan area. The network consists of 4 Oklahoma Mesonet stations and 36 sites mounted on traffic signals. At each traffic signal site, atmospheric conditions are measured and transmitted every minute to a central facility. The Oklahoma Climatological Survey receives the observations, verifies the quality of the data and provides the data to Oklahoma City Micronet partners and customers. One major shortcoming of the system is that it is a much specialised implementation not using open standards or aiming at portability. The same applies to *CORIE* (Center for Coastal and Land-Margin Research, 2009), which is a pilot environmental observation and forecasting system (EOFS) for the Columbia River. It integrates a real-time sensor network, a data management system and advanced numerical models.

Another urban sensing infrastructure is described by Murty et al. (2008). The *CitySense* project uses an urban sensor network to measure environmental parameters and is thus the data source for further data analysis. The project focuses on the development of a city-wide sensing system using an optimised network infrastructure. An important parallel with the work presented in this paper is that *CitySense* also considers the requirements of sensor network setup in an urban environment.

King's College London (2009) presents an urban sensor network for air quality monitoring. The *London Air Quality Network* (LAQN) was formed in 1993 to coordinate and improve air pollution monitoring in London. Increasingly, measurement data are being supplemented by measurements from local authorities surrounding London, thereby providing an overall perspective of air pollution in South East England. Currently, the LAQN consists of about 150 monitoring sites being a very promising approach to urban monitoring as it also offers on-the-fly creation of statistic graphs, time series diagrams and wind plots. However, the network does not make use of open standards as a whole, meaning that it is built up in a proprietary system, although sensor data are accessible over the internet. Despite this shortcoming LAQN can be considered a stand-alone solution of great local significance, though limiting trans-regional inter-linkage with other similar approaches.

Kansal et al. (2007) present the *SenseWeb* project, which attempts to establish a community-driven sensor platform. The project pursues a kind of crowd-sourcing approach by allowing users to import their own sensors in the system. The authors discuss data fusion, transformation issues, and simple geospatial analysis. However, the system is not based on open (geospatial) standards; it just makes use of standard web services.

Paulsen (2008) describes a sensing infrastructure, which aims to combine sensor networks with geospatial visualisation technologies. Thus, the system is highly relevant to this thesis, although it is not deployed in an urban environment. The sensing devices measure rock temperature at ten minute intervals. The project focuses on optimising resource usage, including data aggregation, power consumption, and communication within the sensor network. In its current implementation, the infrastructure does not account for geospatial standards in sensor observations except for the visualisation component.

A geospatial service mashup for environmental data visualisation is implemented in the *nowCOAST* application (National Oceanic and Atmospheric Administration, 2008). *nowCOAST* visualises several types of raw environmental parameters, which

are collected from public data providers. Furthermore, the system offers a 24-hour sea surface temperature interpolation plot.

The major shortcoming of the approaches described above and other related efforts is that their system architectures are at best partly based on open (geospatial) standards.

5.1.2 Sensor Web Characteristics

As mentioned in sub-chapter 3.4, sensor webs have to fulfil several criteria such as interoperability, intelligence and scalability. Sensor webs have only emerged very recently because of increasingly reliable communication technologies, affordable embedded devices and growing importance of sensor data for real-time decision support. This sub-chapter treats the historical development of sensor webs, explains their characteristics and makes a hierarchical distinction between stationary and mobile sensors, as this is a central consideration when implementing sensor webs.

5.1.2.1 Historical Evolution of Sensor Webs

According to Chong and Kumar (2003), the development of sensor webs requires research in three main technological fields, namely the sensors themselves, communication between the sensor units and computation facilities in hardware as well as in software.

The first predecessors of current sensor webs have been developed – as numerous other technologies – for military purposes. During the Cold War, this research resulted for instance in the Sound Surveillance System (SOSUS) to track quiet submarines, or in a network of air defence radars, which were operated in a hierarchical structure and in most cases required human control because of lacking autonomy and intelligence.

The research on more sophisticated technologies such as Distributed Sensor Networks (DSN) started in the early 1980s and was mainly carried out by the Defense Advanced Research Projects Agency (DARPA). At that time, also first ideas in the research area of artificial intelligence (AI) were born. Furthermore, the first operating systems tailored to communication-oriented use on sensor devices, were released (Accent, Mach).

In the 1990s, the idea of network centric warfare (NCW) has been created. In general, this term embodies technology-supported attacks and strategies such as the explicit distinction between weapons and the platforms they are mounted on

allowing for more flexible and controlled assaults. More information about NCW can be found in Alberts et al. (1999).

From the beginning of the current century, sensor developments are coming closer and closer to the original vision of small and cheap sensing devices, which can be rapidly deployed in greatly dynamic environments in an ad-hoc fashion. The main goals of recent research efforts are power-efficient and highly autonomous operation, low latency and self-healing mechanisms. (Chong and Kumar, 2003)

5.1.2.2 Features of Sensor Webs

The criteria for sensor webs are threefold. The first characteristic is *interoperability*, which means that different types of sensors should be able to communicate with each other and produce a common output. The requirement of *scalability* implies that new sensors can be easily added to an existing topology without necessitating aggravating changes in the present hardware and software infrastructure. Finally, *intelligence* means that the sensors are able to think autonomously to a certain degree, which could e.g. result in a data processing ability in order to only send the required data.

Within sensor webs, information is accessed via attribute-based naming instead of explicit addressing, which enables resource-sparing and effective data retrieval. Depending on the capabilities of the sensor, also a history of measured values might be saved resulting in the aptitude to create modular outputs for a certain period of time. This also enables the generation of time series of sensor values originating from a combination of single measurements, which is important for generating motion vectors or mean values over a particular interval to perform plausibility checks.

Furthermore, the Sensor Model Language (SensorML) has been designed for providing an Extensible Markup Language (XML) schema for defining the geometric, dynamic and observational characteristics of a sensor. SensorML and other components of the Sensor Web Enablement (SWE), a conceptual sensor web initiative of the Open Geospatial Consortium (OGC), will be described in detail in sub-chapter 6.2.

These facts make the sensor web a suitable basis for the architecture proposed in this project.

5.1.2.3 Differentiation of Sensors

Figure 8 depicts the decision tree of sensor webs considering different types of sensors and measured data.




Figure 8: Decision Tree of Sensor Web Architectures.

Within this thesis, the differentiation between geo-referenced and non-geo-referenced sensors is at the top of the decision tree, as a meaningful measurement for further geospatial analysis can only be achieved by combining the gauged values with a geographical position. However, a distinction has to be made between mobile and stationary sensors, which is done in the following sub-chapters.

It shall be stated that no differentiation is necessary within this thesis between in-situ sensors, which measure parameters in the direct environment of the sensor, and remote sensors measuring parameters from a certain distance. This is due to the fact that the retrieved data are principally the same from a conceptual viewpoint, i.e. they can both be integrated into pre-defined XML structures such as OGC Observations and Measurements (O&M).

Stationary Sensors

A stationary sensor is deployed at a fixed position where its geographical location is only measured once and need not be retrieved every time measurement data is sent. Nevertheless, the position can be transmitted along with every broadcasted data set in the same manner as with mobile sensors.

One of many existing sample applications for stationary sensors is one of the first deployed sensor webs, namely National Aeronautics and Space Administration's (NASA) sensor network. In the course of this project, sensors have been

established in the Antarctic region to detect environmental changes and report them back to an external control system. (Delin et al., 2003)

Mobile Sensors

In contrast to stationary ones, mobile sensors have to record their current position periodically in order to make measurement results meaningful. Furthermore, an array of previous positions can be established to enable a determination, in which direction the sensor device moves. This requirement involves a dynamic location-aware component within the sensor pod such as a Global Navigation Satellite System (GNSS) module.

Chapter Summary

In the last decade, strongly decreasing sensor costs triggered intense research in sensor technologies and vice versa. This fact, together with miniaturisation efforts and rising awareness of the benefits of automated real-time sensor applications, resulted in the deployment of a number of urban sensor systems. Additionally, the use of live measurements in GIS data analysis systems rapidly gained importance.

Currently, some implementations of geo-sensor networks for urban monitoring applications exist, such as the Oklahoma City Micronet, CitySense, the London Air Quality Network or SenseWeb. The most common deficiency of these approaches establishing urban sensor networks is that their system architectures are at best partly based on open (geospatial) standards limiting interoperability, reusability and flexibility.

5.2 Real-time Data Integration and Sensor Fusion for GIS

As mentioned in the introductory chapter, a key requirement for achieving meaningful geospatial analysis information is high quality of underlying data, where “quality” is defined by the three criteria accuracy, completeness and topicality. Traditionally, the topicality parameter has not received a lot of attention as geospatial data is per definition historic and so far, geospatial analysis has focused on processing static data, with low temporal fluctuations. With the emergence of a variety of real-time data sources, and with the creation of the new paradigms of real-time sensor applications and GIS-motivated situational awareness, topicality of base data is now rapidly gaining significance.

Thus, a sensor fusion mechanism for real-time integration of sensor measurements and other live data has to be created, which combines different data sources in real-time and provides them via standardised interfaces such as OGC Web Feature Service (WFS), Web Map Service (WMS), and Web Coverage Service (WCS). This allows for querying heterogeneous sensor data in real-time and their integration into GIS analysis systems.

To close the gap between real-time data sources and processing technologies, and to overcome the mentioned shortcomings of current approaches, a flexible and portable method has to be created for on-the-fly integration of live sensor data.

5.2.1 Literature Review

Ziegler and Dittrich (2004) mention that the central question in data integration research is how to combine data to offer the user only one view on these data. This requires progress as well in semantics (reduction of losses, common object descriptions etc.) as in data modelling concepts. The authors describe the issues solely from a database perspective.

The same applies to Rittman (2008) who presents Oracle’s approach, which basically consists of a middleware between (web) services and a continuously updated database. Like Sybase Inc. (2008), the Oracle approach is able to detect database events in order to analyse heterogeneous data sources and to trigger actions accordingly. Rahm et al. (2007) developed a more dynamic data fusion approach by suggesting the use of on-the-fly object matching and metadata repositories to create a flexible data integration environment.

These approaches for real-time data integration rely on the costly step of creating a temporal database, which is only suitable in a powerful computation environment without strict performance constraints.

Harrie (2004), and Lehto and Sarjakoski (2005) present web services based on the classic request/response model. Although both methods widely use open geospatial standards, they are not suitable for the integration of real-time data for large volumes of data.

Sarjakoski et al. (2004) introduce a real-time spatial data infrastructure (SDI), which performs several application-specific steps (coordinate transformation, spatial data generalisation, query processing or map rendering and adaptation), but accounts neither for event-based push mechanisms nor for the integration of sensor data.

A number of approaches try to integrate sensor measurements in real-time. However, Srivastava et al. (2006) and Balazinska et al. (2007) present very general concepts in a data integration architecture approach describing the data fusion process chain, but there are no concrete conclusions how the final goal of establishing such an infrastructure could be achieved. A more technical approach addressing measurement data originating from ad-hoc sensor networks is described by Riva and Borcea (2007). The authors discuss challenges to make heterogeneous sensor measurements combinable by creating highly flexible middleware components. The approach is greatly application-motivated and thus very detailed as far as specific implementation details are concerned, but does not offer a generalised model for real-time data integration.

Finally, there are several application-driven data fusion approaches, which try to integrate different kinds of (spatial) data in a very specific way. Different application areas are addressed such as financial data integration (Pan and Viña, 2004), integrating Corporate Performance Management (CPM) and Business Intelligence (BI), clinical data integration (Choi et al., 2005 and Yoo et al., 2003) or heterogeneous data visualisation (Kansal et al., 2007). However, these approaches do not focus on the use of open standards and they are not considering the combination of push mechanisms and real-time data integration.

5.2.2 Shortcomings of Previous Approaches

Real-time data integration and data fusion are non-trivial research areas. As data fusion originated in the domain of computer science, very few approaches exist, which are dedicatedly designed for geospatial data. Thus, integration of sensor data into GIS analysis systems at the moment mostly happens via the laborious interim step of a temporary physical database. This is not desirable in an automated GIS workflow as the database can easily become a bottleneck for large data volumes. Moreover, this approach unnecessarily adds another component to the overall

workflow, which can result in substantial losings in performance. Thus, a need arises for an approach towards on-the-fly integration of sensor measurements.

Furthermore, current data integration systems are often inflexible for the adaptation to other data structures, while being characterised by a high degree of external dependence to particular data fusion software.

This is often the result of lacking compliance of the software with open data provision standards restricting response formats to very application-specific and proprietary data structures. This is a central issue particularly in integrating real-time sensor data into GIS products, which use a variety of well-established data structures, but as well outside the geospatial domain. Apart from that, output data in standardised formats can be easily integrated into mapping APIs such as Google Maps or Microsoft Virtual Earth, which have recently gained extensively in popularity.

Chapter Summary

Traditionally, topicality as a quality parameter for geospatial data has not received a lot of attention as geospatial data is per definition historic and so far, geospatial analysis has focused on processing static data, with low temporal fluctuations. With the emergence of a variety of live data sources, topicality of base data is now rapidly gaining significance calling for efficient and flexible real-time data integration mechanisms.

Currently, integration of sensor data into GIS analysis systems mostly happens via the laborious non-desirable interim step of a temporary physical database, which limits performance and increases the probability of component failure.

Furthermore, current data integration mechanisms are often inflexible for the adaptation to other data structures, while being characterised by a high degree of external dependence to particular data fusion software.

5.3 Web-based Geospatial Data Analysis

Traditionally, GIS systems in general and geo-processing in particular have been implemented in monolithic solutions with tight coupling of data, processing and visualisation in proprietary systems. Currently, there is a shift occurring towards web-based geo-processing, which can be considered the server-side execution of algorithms that operate on geospatial data.

5.3.1 The OGC WPS Standard

However, during the last decade, Service Oriented Architecture (SOA) principles are increasingly being applied for the realisation of open, distributed spatial data infrastructures (SDI). The Open Geospatial Consortium (OGC) has created a set of web service interface standards for publishing and accessing image maps, features, coverages and sensor data, as well as their geographic metadata. These specifications are now quite mature and constitute the basis of most web-enabled geospatial systems that are currently being developed.

Following this service-oriented trend, the OGC started an attempt to also integrate web-processing functionality in a workflow-oriented infrastructure by creating the Web Processing Service (WPS) standard (Schut, 2007). WPS defines an open interface that enables publishing of geospatial processes and allows for the development of software clients that discover and bind to those processes.

5.3.2 Literature Review – WPS Implementations

Currently, several organisations and open-source projects provide standards-compliant implementations of the WPS 1.0 specification:

- The 52°North open-source project has developed a fully Java-based implementation of the WPS 1.0.0 specification. The software is provided as an open framework in order to allow for other processing modules to be added. In addition to the server, 52°North also supplies a WPS client plug-in for the open-source UDIG desktop GIS software. (52North, 2009)
- PyWPS is a server implementation of the WPS 1.0.0 specification written in Python with native support for GRASS GIS. (Cepicky, 2009)
- WPSInt is a Java-based implementation of the outdated WPS 0.4.0 specification that is built using the Java Spring framework. It comprises a Java server module, which handles WPS requests and responses, and a generic WPS client. (CollabNet Inc., 2009a)

- Deegree's WPS implementation is fully Java-based and provides an implementation of a Java framework for the development of Web Processing Services. (lat/lon GmbH, 2009)
- PLWPS is an open-source implementation of the WPS 1.0.0 specification realised in ASP.Net including a basic web client application written in PHP (Hypertext Preprocessor). (SourceForge, 2009)
- A WPS add-on for the GeoServer open-source project is currently being developed. It aims to incorporate the Sextante GIS processing modules. (OpenGeo, 2009)

In parallel, OGC has been working on a Web Coverage Processing Service (WCPS) that defines a protocol-independent language for the extraction, processing, and analysis of multi-dimensional gridded coverages representing sensor, image, or statistical data. Services implementing this language provide access to original or derived sets of geospatial coverage information, in forms that are useful for client-side rendering, input into scientific models, and other client applications (Baumann, 2009)

5.3.3 WPS Profiles

The WPS specification (Schut, 2007) defines a generic mechanism that can be used to describe and publish any sort of geospatial process. However, to achieve far-reaching interoperability, each process has to be specified in a separate document, which can be considered an application profile of the WPS specification containing the following parameters.

- *URN*: unique identification of the process (mandatory)
- *Reference response*: reference to a DescribeProcess request for the specific process (mandatory)
- *Description*: a human-readable document that describes the process and its implementation (optional, but recommended)
- *WSDL description*: a WSDL document describing this process (optional)

The creation of a WPS application profile is not trivial and straight-forward. First, the definition of a unique URN (Uniform Resource Name) is problematic as the OGC has not defined consistent rules for specifying URNs. Secondly, there is no clear definition of how these profiles should be hierarchically structures. Nash (2008) presents an approach, in which a distinction is made between generic processes providing essential GIS functionality, and specialist processes. In a further step, the

author suggests the differentiation between processes that operate on raster data from those working on vector data. This is due to the existing gap between vector and raster data processing in today's GIS community.

5.3.4 Asynchronous processing

An essential aspect when dealing web-based geo-processing is that geospatial computations require a certain amount of time to be executed. For some operations processing large amounts of data in a complex algorithm, this calculation process may take longer than the typical Hypertext Transfer Protocol (HTTP) time-out duration. Thus, WPS cannot only operate in the classical request/response model, but has to account for asynchronous processing as well, providing the possibility to send the computation result outside the HTTP response. This asynchronous communication can happen in two ways after the client has sent a WPS "Execute" request to the server:

- the client regularly polls the server to check whether the request is finished
- the client gets notified by the server when the process has completed

The current WPS version (1.0.0) defines the polling approach as the basis for its asynchronous operations. Hence, after the client has sent the "Execute" request to the server, an execute response message is immediately returned including a URL indicating the status of the process. After completion of the process, the WPS service updates the document and provides the client with the URN, at which it will publish the response document. One of the disadvantages of this polling method is the significant overhead in messages that are exchanged as the client needs to continuously poll the server to find out whether a process has been completed.

Thus, the approach of notifying the client if the process has been completed seems to be a more suitable and optimised approach considering the exchange of messages between client and server. However, this method also requires a set of underlying messaging standards for notification etc. The World Wide Web Consortium (W3C) provides a number of relevant standards (WS-Notification, WS-Addressing etc.), but they are not fully acknowledged and integrated in current OGC specifications. Furthermore, these W3C services require the use of the Simple Object Access Protocol (SOAP) as a communication protocol, which is only one of several possibilities when using OGC WPS.

Chapter Summary

Traditionally, GIS systems in general and geo-processing in particular have been implemented in monolithic solutions with tight coupling of data, processing and visualisation in proprietary systems. Currently, there is a shift occurring towards web-based geo-processing, which can be considered the server-side execution of algorithms that operate on geospatial data.

Following the service-oriented trend of the past decade, the Open Geospatial Consortium (OGC) started an attempt to integrate web-processing functionality in a workflow-oriented infrastructure by creating the Web Processing Service (WPS) standard.

Research areas, which are still in early and immature phases, comprise the definition of web-based processing semantics, WPS profiles for specific use case domains, and asynchronous processing of real-time data.

LIVE GEOGRAPHY APPROACH

LIVE GEOGRAPHY APPROACH

6 "Live Geography" Approach

Despite projections that cities would die as they are barely artefacts from the industrial age, urban population has reached a historic peak. In 2008, for the first time in history, more than half the world's population lives in cities; tendency rising. This fact, coupled with a far-reaching research aspect of multi-layered complexity in urban processes, yielded increased interest in exploring city dynamics. This is especially true for real-time monitoring as enabling technologies have become technically feasible and affordable for pervasive deployment, as elucidated in the introductory chapter.

This chapter illustrates the Live Geography approach proposing a technical infrastructure enabling a wide variety of urban monitoring applications. It comprises an illustration of technological background topics (standardisation, disambiguations of geo-workflow components etc.), as well as a description of the key requirements of the architectural modules, i.e. geo-sensor webs, real-time data integration and sensor fusion, location-aware event processing and alerting, and web-based geo-processing. Possible application areas as well as the broader impact of the Live Geography developments on the urban environment are elucidated in chapters 8 and 9, respectively.

6.1 Background

A key requirement for performing meaningful geospatial analysis processes is high quality of underlying data. As mentioned in the introductory chapter, quality criteria for geospatial data comprise accuracy, completeness and topicality. The topicality parameter is still neglected in many cases, while the first two parameters have recently received a lot of attention.

One possible reason for this is that geospatial data is per definition historic and so far, GIS-based analysis has focused on processing static data, with low temporal fluctuations. With the emergence of a variety of real-time data sources particularly in cities (e.g. sensor networks, RFID-based systems, traffic counters etc.), the creation of the new paradigm of GIS-based situational awareness, topicality of base data is now rapidly gaining significance.

Utilisation of urban real-time data in GIS applications requires a rethinking process of existing practices. Currently, geospatial data analysis is performed in costly specialised software involving a high degree of manual intervention for data gathering, pre-processing and quality assurance. Furthermore, geospatial analysis is per definition applied on historic data with long cycles from data generation to

analysis output and a potential real-world impact. As the Live Geography approach revolves around geo-sensor networks it aims to create a real-time feedback loop between sensor measurements, data processing, and evaluation and usage of the resulting information for supporting short-term decisions in near real-time. The basic architecture for such applications is illustrated in Figure 9.




Figure 9: Basic Architecture Components of the Live Geography Infrastructure.

Generally speaking, the infrastructure shown in Figure 9 can be sub-divided in five components, i.e. stand-alone parts, which have to be conflated. Component **1** is the geo-sensor network itself including measurement devices Global Navigation Satellite System (GNSS) connectivity and basic processing capabilities. Component **2** covers the communication of the sensor network with a data centre, including data transmission, harmonisation and integration. Component **3** comprises the application-specific analysis of the sensor data on the server side. This operation does not only include pure sensor data processing, but also the integration of static and legacy geospatial data. Component **4** treats the presentation of analysed data depending on the particular requirements of end users or user groups. Component **5** finally deals with the archiving process of sensor data in order to store longer-term measurement histories and to disburden the sensor network itself, which can only store a limited amount of data.

6.2 Standardised System Architecture

To automate these procedures, a service-oriented approach supported by the extensive use of open standards has to be pursued. In the construction of the technical infrastructure, different design principles (Service Oriented Architectures – SOA, modular software infrastructures etc.) were taken into account to ensure flexibility, reusability and portability of the components and the overall infrastructure. Figure 10 shows the modular architecture and the interfaces between the different components.




Figure 10: Live Geography – Workflow Illustration with Interfaces of the Architecture Approach.

Following principles of SOA and sustainable infrastructure development, a use case independent data collection, processing and information provision architecture has been conceived within this thesis. The infrastructure covers the whole process chain from sensor network development via measurement integration, data analysis and information visualisation, as presented in Figure 10. Thus, the Live Geography approach builds the architectural bridge between domain-independent sensor network development and use case specific requirements for end user tailored information output.

The first module in Figure 10 illustrates the sensor network including a description of the sensor and the containing platforms. This is achieved by an OGC Sensor Model Language (SensorML) document including different parameters about the sensing platform itself such as contained sensors, orientation, location, supported measurement methods, as well as input and output parameters of the mounted sensors. At the same time, the sensor can be directly accessed via the Sensor Planning Service (SPS) interface to adjust sensor parameters remotely. In practice, this would e.g. mean that the measurement cycle is changed from hourly to once a minute if the system predicts a high chance of a flood water event occurring in the near future.

In a second step, raw sensor measurements are harmonised by encapsulating them into pre-defined XML-based OGC Observations and Measurements (O&M) format.

O&M defines an abstract data model (XML schema), which allows for the provision of sensor measurements (numerical values, raster images, binary states, complex or combined measurement data etc.) in a structured and standardised format. These O&M-encoded sensor data are then served via the OGC Sensor Observation Service (SOS) interface. Sub-chapter 6.4 describes the sensor device including its software infrastructure in more detail.

This point in the workflow is also where usually Complex Event Processing (CEP) mechanisms are integrated. This means that e.g. averaged values, certain measurement parameters or complex patterns in sensor data trigger so-called "events", as described in sub-chapter 6.6. A practical example would be that the exceedance of an air quality threshold triggers appropriate geospatial analysis functions (interpolation, spatio-temporal correlation with other parameters such as traffic data etc.) to process raw sensor data and create meaningful and intuitively assessable information layers. Furthermore, CEP can serve for assuring quality control in measurement data. This happens by different kinds of plausibility checks in the spatial and temporal domains, such as time line trend comparison, outlier detection, spatial averaging etc.

Next, the O&M-encoded measurements are harmonised by the sensor fusion component, and provided via standardised OGC Web Feature Service (WFS), Web Map Service (WMS) and Web Coverage Service (WCS) interfaces. In fact, the real-time data integration module acts as an on-the-fly "translator" from the O&M format to the WFS interface. Technically, the component interprets O&M-encoded data and can then autonomously convert this input to other output formats. To enable this conversion, different parameters such as coordinate reference system, measurement parameters, mathematical units, data formats etc. have to be conveyed to the sensor fusion module, which can then perform an automatic real-time transformation. More details about the sensor fusion component are illustrated in sub-chapter 6.5.

As it can be seen from Figure 10, the creation of this standards-based infrastructure naturally has to account for several open standards such as the concepts within the Sensor Web Enablement (SWE), the Geography Markup Language (GML), WFS or WMS, as described in the following sub-chapters.

6.2.1 Web Service Standards

Generally speaking, when setting up interoperable web services, the application has to be compliant to the Hypertext Transfer Protocol (HTTP), the Extensible Markup Language (XML) as well as the Simple Object Access Protocol (SOAP) for data

exchange and finally to the Web Service Description Language (WSDL) and the Universal Description, Discovery and Integration (UDDI) for publishing a service based on the data provided by the sensor web. Moreover, the application should possibly allow the integration of the Business Process Execution Language (BPEL), which serves for modulating business processes as web service chains and editing them for coordinated collaboration with business partners. The relationship between these standards is illustrated in Figure 11.




Figure 11: Web Service Standards Structure. (IBM Corporation, 2005)

Figure 11 shows the well-known "publish-find-bind" web service triangle. It basically says that a service provider can publish its web services in a Universal Description, Discovery and Integration (UDDI) registry containing a description of the service in the standardised Web Service Description Language (WSDL). A client can retrieve this service description from the registry and send a request to the service instance itself using the according parameters as advertised in the WSDL document. This process is well-established in service-oriented environments and is thus not described in more detail at this point.

6.2.2 Sensor Web Enablement Initiative

For standardisation in the field of geo-sensor networks, the Open Geospatial Consortium (OGC) brought the Sensor Web Enablement (SWE) initiative into being, whose goal is to make different types of sensing devices discoverable, accessible and possibly controllable over the web, as shown in Figure 12.

SWE comprises seven OGC standards or interoperability programme reports, the Sensor Model Language (SensorML), Observations and Measurements (O&M), the Sensor Observation Service (SOS), the Sensor Alert Service (SAS), the Transducer Model Language (TransducerML), the Sensor Planning Service (SPS) and the Web Notification Service (WNS), which are described in the following sub-chapters. The overall goal of these standards is to enable the creation of interoperable and scalable sensor networks, which are built up in a service-oriented fashion. (Botts et al., 2007)




Figure 12: General SWE Idea and Architecture. (Botts et al., 2007)

6.2.2.1 Sensor Model Language

The Sensor Model Language (SensorML) is a general schema for describing functional models of the sensor. It provides an Extensible Markup Language (XML) schema for defining the geometric, dynamic and observational properties of a sensor. Thus, SensorML serves for discovering different types of sensors, supporting processing and analysis of the retrieved data as well as the geo-location of observed values. Information provided by SensorML includes observation and geometry characteristics, as well as a description and a documentation of the sensor, and a history of the component's creation, modification, inspection or deployment. (Botts, 2007)

SensorML aims to provide descriptions of sensors and sensor systems for inventory management, to supply sensor and process information in support of resource and

observation discovery. Furthermore, SensorML aims to allow for geo-locating observed values (measured data), to provide performance characteristics (e.g., accuracy, threshold, etc.) and to offer an explicit description of the process, by which an observation was obtained.

Figure 13 shows the element hierarchy of SensorML, whereby the root element is always named *Sensor* and the nine possible child elements are used to identify the sensor and to describe its capabilities and characteristics. SensorML encodes all processes and components as application schema of the Feature model in the Geographic Markup Language (GML) Version 3.1.1.




Figure 13: Element Hierarchy in SensorML. (Botts, 2004)

An important capability of SensorML that has to be mentioned considering the background of this thesis is the formation of sensor classes, i.e. sensing devices with the same properties. In SensorML, a sensor array defines a set of devices of the same type at different locations, whereby a sensor group describes several sensors that operate together to provide one collective observation.

6.2.2.2 Observations and Measurements

Observations and Measurements (O&M) is the pendant to SensorML in the field of the actual study of phenomena. That is, it provides a description of sensor observations and measurements in the form of general models and XML encodings. The term *observation* is defined as "an action with a result which has a value describing some phenomenon". The O&M standard labels several terms for the measurements themselves as well as for the relationship between them, whereby the extent is limited to measurement results, which are expressed as quantities, categories, temporal or geometrical values as well as arrays or composites of these.

Naturally, the monitored values normally require a reference system to enable a meaningful interpretation and further processing by providing a context for the results. (Cox, 2007)

It shall be stated that despite OGC's tendency towards using the geographical position as the central and connecting element in geospatial standards, the location parameter is considered a regular measurement within O&M. This means that the position is equipollent with other measurands such as time, air temperature or satellite images.

6.2.2.3 Sensor Observation Service

The Sensor Observation Service (SOS) provides a service to retrieve measurement results from a sensor or a sensor network. In other words, the SOS groups a collection of possibly heterogeneous sensors, as illustrated in Figure 14, and provides their measurements via a standardised service interface.




Figure 14: General SOS Usage Concept. (Adapted from Na and Priest, 2007)

The SOS specification defines the operations offered by a specific sensor, whereat the minimum collection of methods comprises *GetCapabilities*, *DescribeSensor* and *GetObservation*, which return information about the observations and measurements supported by the SOS. The kinds of data provided by a sensor and the sensor types themselves can be fetched from a sensor registry. SOS references the O&M specification for encoding sensor observations, and the TransducerML and SensorML specifications for modelling sensors and sensor systems. (Na and Priest, 2007)

6.2.2.4 Sensor Alert Service

The OGC Sensor Alert Service (SAS) specifies interfaces (not a service in the traditional sense) enabling sensors to advertise and publish alerts including according metadata. SAS can be used by clients to subscribe for sensor data alerts with some spatial and property-based constraints. Also, sensors can be advertised to the SAS allowing clients to subscribe for the sensor data via the SAS, as shown in Figure 15. SAS, which is currently in its version 0.9.0, is not released as an official OGC standard. (Simonis, 2007a)




Figure 15: General SAS Functionality. (Simonis, 2007a)

In the SAS context, "alerts" are not only understood in the classic meaning of the word, i.e. an automatic signal or notification indicating that an event has fired (e.g. a message in case of threshold exceedance), but in a broader context. Alerts are defined as "data" sent from the SAS to the client, which may as well comprise alerts/notifications (e.g. OGC Web Notification Service [WNS]) as observational data (measurements matching pre-defined criteria).

As Figure 15 illustrates, SAS uses the Extensible Messaging and Presence Protocol (XMPP) for the delivery of sensor notifications. Thus, SAS leverages an XMPP server, which can be embedded directly in the SAS or act as a separate service. SAS notifications are provided via a Multi User Chat (MUC) for each registered sensor. To receive notifications, a client has to join the specific MUC.

6.2.2.5 Sensor Planning Service

The Sensor Planning Service (SPS) provides a standardised interface for collection assets and aims at automating complex information flows in large networks. This requires the support of various capabilities as well as different request processing systems, as described below.

The OGC recommendation defines interfaces for requesting information containing the capabilities of a SPS, for retrieving the practicability of a sensor planning request, for determining the status of such a request as well as for submitting, updating or cancelling a sensor planning request. (Simonis, 2007b)

Thus, SPS can be seen as the implementation of a bi-directional communication between a client and the sensor network. The client can send instructions to the network, which can be triggered by different kinds of events. One possibility is that a system administrator instructs the sensor network to send measurement values every five minutes instead of every hour in case a water level threshold is exceeded. Another commonly used sample functionality is the control of a remote camera to change orientation or focus.

6.2.2.6 Web Notification Service

The last part of the SWE model is the Web Notification Service (WNS), which is a service enabling a client to perform asynchronous dialogues, that is message exchanges, with other services. This process is especially expedient when several services are required to comply with a client's request, or when an according response is only possible under considerable delays.

Principally, the service comprises two kinds of communications, the one-way-communication and the two-way-communication. The former, also known as simple WNS, for which users have to be registered, can be realised via the transmission of messages over e-mail, HTTP, Short Message Service (SMS), instant messaging, telephone, letter or fax. The latter system, viz extended WNS, is able to receive user notification responses. (Simonis and Wytzisk, 2003)

6.2.2.7 SWE Common

An important ongoing effort in SWE development is the establishment of a central SWE Common specification. Its goal is to optimise common usage and maximise reusability by grouping common elements for several standards under one central specification. SWE Common is not directly part of the SWE initiative, but tries to reduce redundancy between the seven standards. In other words, the SWE

Common namespace specification aims at collecting elements, which are used in more than one standard of the SWE family. SWE Common will mainly comprise very general elements such as counts, quantities, time elements or simple generic data representations.

6.2.2.8 Sensor Web Registries

Furthermore, sensor web registries play an important role in sensor network infrastructures. However, they are not decidedly part of SWE, as the legacy OGC Catalogue Service for Web (CSW) is used.

The registry serves to maintain metadata about sensors and their observations. In short, it contains information including sensor location, which phenomena they measure, and whether they are static or mobile. Currently, the OGC is pursuing a harmonisation approach to integrate the existing CSW into SWE by building profiles in ebRIM/ebXML (e-business Registry Information Model).

6.2.2.9 Sensor Web Enablement – Summary

Summarising, SWE aims at discovery and query of sensor systems, observations, and observation procedures over the internet. This process comprises determination of a sensors' capabilities and quality of measurements, access to sensor parameters that automatically allow software to process and geo-locate observations, retrieval of real-time or time-series observations and coverages in standard encodings, tasking of sensors to acquire observations of interest, and subscription to and publishing of alerts to be issued by sensors or sensor services based upon pre-defined criteria. Figure 16 shows the functional connection of the SWE standards.




Figure 16: Functional Connection of SWE Components. (Adapted from Botts et al., 2007)

The OGC is developing SWE in tight coordination with other geospatial standards for data representation, provision and processing. This enables simple integration of sensor networks into existing Spatial Data Infrastructures (SDI) and data analysis systems by providing standardised service interfaces and API (Advanced Programming Interfaces).

6.2.3 Other Relevant OGC Standards

Apart from these sensor-related standards, other data provision services and data encoding formats specified by the OGC are relevant to the Live Geography approach.

6.2.3.1 Web Feature Service

The OGC Web Feature Service (WFS) specification, defines a service-oriented interface for data access and manipulation operations on geographic features using HTTP as the distributed computing platform. Via these interfaces, a web user or service can combine, use and manage geo-data. (Vretanos, 2005)

Basically, WFS enables access to raw feature data (in contrast to the Web Map Service) and provides a variety of output formats (Geographic Markup Language [GML], geoRSS, Keyhole Markup Language [KML], amongst others) facilitating the integration of feature data into Geographic Information Systems.

The Web Feature Service allows for following manipulations on geographic features:

- Create a new feature instance
- Delete a feature instance
- Update a feature instance
- Lock a feature instance
- Get or query features based on spatial and non-spatial constraints

Thus, the WFS operations support INSERT, UPDATE, DELETE, LOCK, QUERY and DISCOVERY transactional mechanisms using HTTP as the distributed computing platform protocol.

The Web Feature Service protocol follows the common OGC communication sequence by defining mandatory operations, as illustrated in Figure 17. The *GetCapabilities* operation returns a document indicating the service's offerings, i.e. which data layers are available, which versions are supported, which output formats are provided etc. The *DescribeFeatureType* method returns a describing XML schema for a specific feature type. This information can then be used to send the actual requests for data retrieval (*GetFeature*) and manipulation (*Transaction*).




Figure 17: WFS Protocol Diagram. (Computational Geoscience Group, 2009)

6.2.3.2 Web Map Service

Similarly to the Web Feature Service, the Web Map Service (WMS) provides a service interface for accessing geographic data. However, while WFS output are actual geographic feature data, WMS produces a map from geographic data. Output formats comprise Portable Network Graphics (PNG), Graphics Interchange Format (GIF) or Joint Photographic Experts Group (JPEG) images, or vector-based elements such as Scalable Vector Graphics (SVG).

WMS also follows the common OGC request/response model by offering the mandatory *GetCapabilities* and *GetMap* operations. More about WMS data retrieval, protocol structures and supported operations can be found in de la Beaujardiere (2006).

6.2.3.3 Web Coverage Service

Like WFS and WMS, the OGC Web Coverage Service (WCS) enables access to geographic features via the common protocol model. WCS provides "coverages", i.e. "digital geospatial information representing space-varying phenomena". Data are served along with their detailed descriptions, a comprehensive syntax for requests for these data, and their original semantics, which may be interpreted, extrapolated, etc. In other words, WCS returns coverages containing space-varying phenomena that represent a spatio-temporal domain to a range of properties. WCS offers three mandatory operations for data access, namely *GetCapabilities*, *DescribeCoverage*, and *GetCoverage*, which are conceived in the same way to the described WFS and WMS methods (Whiteside and Evans, 2008)

6.2.3.4 Geography Markup Language

The Geography Markup Language (GML) consists of an XML schema, which serves for modelling, exchanging and storing geographic data. GML provides a set of objects for the description of coordinate reference systems, geographic elements such as points, lines, topologies, units and many more.

Originally, GML has been created regarding following design principles:

- Creation of a vendor-independent framework for the exchange of geographic data
- Providing a basis for the storage and the porting of application schemas and data sets
- Supporting domain-specific schema descriptions

GML Version 3 Properties

The current third version of GML originated from several shortcomings of GML 2 such as the restriction to simple geometric structures defined by two coordinates. For the creation of more complex structures, a set of features³ describes an object in the new version.

GML 3 comprises following enhancements versus its predecessor:

- Description of more complex structures like non-linear three-dimensional (3D) objects or dynamic features including temporal properties

³ According to International Organization for Standardization (2005a), a feature is an "abstraction of real-world phenomena".

- Support of spatial and temporal reference systems and units
- Standardised visualisation of features

In order to satisfy these improvements, the set of basic schemas had to be expanded by the factor eight. Despite these extensions, GML 3 is wholly compatible with former versions until GML 2.1.2. (Portele, 2007)

Class Hierarchy

Figure 18, which was obtained from International Organization for Standardization (2005a), illustrates the class hierarchy of GML 3.1.0 in a simplified Unified Markup Language (UML) structure.




Figure 18: Class Hierarchy of GML 3.1.0. (International Organization for Standardization, 2005a)

The elements beginning with the word *Abstract* represent objects of their class, e.g. *AbstractTopology* stands for any GML topology. The normative schemas are arranged around these objects.

Schemas

The GML structure is based on numerous schemas, which enable the visualisation of different types of geographical data.

The World Wide Web Consortium (W3C) XML schema part 2 (Biron und Malhotra, 2004) defines various "simple" types, with which values can be illustrated as letters without any extra markup. As GML is made up by XML schemas, these simple types are used as far as possible. Nevertheless, if the W3C schema does not contain appropriate types, GML base types can be derived, which are described in Annex C of Portele (2007).

The basic element on top of the class hierarchy is termed *AbstractObject*, as it can be seen from Figure 18. This element does not comprise a defined type and is thus an XML schema named *anyType*. For the creation of an identifiable GML object, the *AbstractGML* element, which contains all other GML elements, is placed at the disposal. The combination of this element with the property *gml:AbstractGMLType* forms a frequently used pattern, whereby each GML object is epitomised by a global element declaration, which has an associated XML schema type definition. The sub-elements shown in Figure 18 are described in detail in International Organization for Standardization (2005a).

The gmlBase Schema

The *gmlBase* schema defines components, which form the GML model:

- A basic XML type, from which all GML elements can be derived
- A pattern and components for GML properties
- Patterns and components for field arrays
- Components for the combination of meta data with GML objects
- Components for the creation of definitions and dictionaries

GML Properties

A property in GML describes a characteristic of an object. A GML element can only be a GML property if it is a child element of the root element *AbstractObject*. These properties can be described by two methods; the property can either be defined by

the content of the element, or it is available at another location, which is referenced by an *xlink:href* attribute in the property element.

Additional GML Schemas

Apart from the basic schemas, several additional ones are defined:

- *Features*: meaningful objects, which are classified in numerous domains to describe real-world objects such as road, river, person etc.
- *Basic geometry*: geometric primitives (not decomposable into other elements), aggregates and complex elements
- *More geometric primitives*: additional geographic simple types to describe real-world phenomena requiring a more complex model
- *Geometric complex, composites and aggregates*:
 - Geometric complexes: closed combinations of geometry disjoint primitives in a common coordinate reference system (CRS)
 - Geometric composite: geometric complex with underlying core geometry having the same form as a primitive
 - Geometric aggregates: compositions of geometry elements; do not have an additional internal structure
- *Coordinate Reference Systems*: description of geodetic and non-geodetic Coordinate Reference Systems (CRS)
- *Topology*: description of the characteristics of an object, which are firm under constant deformation, e.g. an ellipse is topologically the same as a circle as transformation of one into the other occurs by simple stretching
- *Temporal information and dynamic features*: description of temporal geometry, topology, reference systems and properties of geographic data including the characterisation of moving objects
- *Units, measures and values*: definition of quantitative values referenced by a scale or units of measure
- *Directions*: description of an object's direction, orientation, heading etc. by a standard property element
- *Observations*: modelling of the observation process of different devices, instruments or persons
- *Coverages*: two possible ways of representation:

- A set of discrete location-value pairs
- A description of the spatio-temporal domain and a description of the set of values

The schemas listed above are described in detail in International Organization for Standardization (2005a).

6.3 Geo-workflow Disambiguation – Sensor Fusion, Data Fusion and Geo-processing

In the creation of the Live Geography approach, a clarification of the terms "data fusion", "sensor fusion" and "geo-processing" has to be performed in order to define clear research terminology. The originating question is whether a clear distinction between sensor fusion, data fusion and geo-processing can be made; and if yes, how do they relate to each other within the overall processing workflow.

6.3.1 Background – Literature

Historically speaking, the term "data fusion" has originated in the field of computer sciences and has then been adapted for specialised domains such as earth observation, remote sensing, data mining or pattern recognition. Furthermore, its meaning has been transferred to other domains like bioinformatics, business intelligence or intelligent transportation systems causing a wide spread of different definitions and usage forms. On the contrary, the term "geo-processing" has been coined by a more homogeneous "geo"-community.

The following table summarises two of the most common and far-reaching for data fusion and geo-processing definitions:

Data-Fusion.org (Ecole des Mines de Paris, 2009)	ISO 19119 Standard (International Organization for Standardization, 2005b)
<i>Data Fusion definition:</i> ■ Data fusion exploits the synergy offered by the information originating from various sources. ■ Data fusion allows formalising the combination of these measurements, as well as to monitor the quality of information in the course of the fusion process. ■ It means an approach to information extraction spontaneously adopted in several domains before this was expressed as "data fusion". This approach is based upon the exploitation of the synergy offered by the information originating from various sources. ■ Fusion may operate at any of the three semantic levels: measurements (fusion of measurements), attributes (fusion of attributes) and rules (fusion of decisions or rules), with possible crossings between levels: ■ <i>Fusion of attributes</i> consists in merging the attributes of a same object, derived from two representations. ■ <i>Fusion of analysis</i> consists in aggregating representations into a new representation. ■ <i>Fusion of representations</i> is defining and performing meta-operations applicable to representations to obtain a new representation. <i>General remark:</i> ■ the term processing is not used in data fusion definitions	<i>Data Fusion definition:</i> ■ Geographic data fusion is a framework of services for the synthesis of data from different sources. It aims to obtain information of higher quality. ■ Geographic fusion is accomplished through organising, relating and linking disparate sources of location-based information. ■ Geographic fusion relates to the fusion of two or more geographic elements and the fusion of geographic elements with other sources of structured and unstructured data from distributed sources. <i>Geo-processing definition:</i> ■ Processing services are services that perform large-scale computations involving substantial amounts of data. ■ Processing services partially modify the attributes of features. ■ A processing service does not include capabilities for providing persistent storage of data or transfer of data over networks. ■ Geographic processing services taxonomy ■ Geographic processing services – spatial ■ Geographic processing services – thematic ■ Geographic processing services – temporal ■ Geographic processing services – metadata <i>General remark:</i> ■ this standard makes a clear distinction between data fusion and geo-processing

Table 1: Definitions for Data Fusion and Geo-processing

6.3.2 Data Fusion Definitions

In many recent definitions, data fusion is explained using the concrete example of (multi-sensor) data fusion – often referred to as *sensor fusion* – a way of combining measurements of physical quantities from one or more sensor sources such that the resulting information is more accurate, more reliable or more complete.

Llinas et al. (2004) characterise data fusion using four levels of data fusion. To categorise the different types of fusion processes as Data Fusion model was developed by the Joint Directors of Laboratories (JDL) Data Fusion Group. The four data fusion levels are illustrated in Figure 19.

- Level 0: Signal Assessment – estimation and prediction of signal states
- Level 1: Entity Assessment – estimation and prediction of entity parametric and attributive states (i.e. of entities considered as individuals)
- Level 2: Situation Assessment – estimation and prediction of the structures of parts of reality (i.e. of relations among entities and their implications for the states of the related entities)
- Level 3: Impact Assessment – estimation and prediction of the utility/cost of signal, entity or situation states – including predicted impacts given a system's alternative courses of action




Figure 19: JDL Data Fusion Levels. (Steinberg and Bowman, 2004)

It is not necessary to process these levels in a given order; usually it has been avoided to treat them hierarchically. Nonetheless, there is often a natural progression from raw measurements to entity states, to situation relationships, and to utility prediction, in which data are combined as suggested by the level ordering. Such composition of estimated signals, entities, and aggregates per Levels 0-2 is

quite natural (Steinberg & Bowman, 2004). The progression from raw measurements to entities (Data Fusion Level 0-1) therefore can be referred to as *Sensor Fusion*.

In contrast to JDL, Ecole des Mines de Paris (2004) lists various data fusion characteristics, as shown in Table 1. This implies that fusion may operate at any of the three semantic levels: measurements (fusion of measurements), attributes (fusion of attributes) and rules (fusion of decision or rules), with possible crossings between levels.

According to Waltz and Llinas (1990), there are many advantages in using data fusion including robustness and reliability (the system remains operational even if one or several sources of information are missing or malfunctioning), extended coverage in space and time, increased dimensionality of the data space, increased the quality of the deduced information and reduced ambiguity, i.e. more complete information provides better discrimination between available hypotheses.

However, the ISO19119 standard on *Geographic Information – Services* (International Organization for Standardization, 2005b) states that "Geographic data fusion is a framework of services for the synthesis of data originating from different sources. It aims at obtaining information of greater quality; the exact definition of greater quality will depend upon the application. Geographic fusion is accomplished through organising, relating and linking disparate sources of location-based information. Geographic fusion relates to the fusion of two or more geographic elements and the fusion of geographic elements with other sources of structured and unstructured data from distributed sources."

6.3.3 Geo-processing Definitions

The term "geo-processing" has also been defined by a number of institutions. According to the ISO19119 standard (International Organization for Standardization, 2005b), processing services are "services that perform large-scale computations involving substantial amounts of data. [...] A processing service does not include capabilities for providing persistent storage of data or transfer of data over networks." Moreover, the standard specifies service taxonomies to differentiate between different kinds of processing services, as listed in Table 1.

The OGC WPS 1.0.0 standard (Schut, 2007) characterises a web processing service in a more technical and service-oriented way. "A WPS can be configured to offer any sort of GIS functionality to clients across a network, including access to pre-programmed calculations and/or computation models that operate on spatially

referenced data. [...] This Web Processing Service is targeted at processing both vector and raster data."

ESRI (2009) defines the general term "geo-processing" using highly geography-motivated and method-oriented terminology. "Geo-processing is based on a framework of data transformation. A typical geo-processing tool performs an operation on a dataset (such as a feature class, raster, or table) and produces a new dataset as the result of the tool. Each geo-processing tool performs a small yet essential operation on geographic data, such as projecting a dataset from one map projection to another, adding a field to a table, or creating a buffer zone around features." Figure 20 shows the general geo-processing scheme as defined by ESRI.




Figure 20: ESRI Geo-processing Scheme. (ESRI, 2009)

Conclusion

In conclusion, it can be stated that *data fusion* and *geo-processing* overlap to a certain degree in the sense that both combine different kinds of data sources to achieve a meaningful outcome by applying various types of algorithms.

It is striking that data fusion definitions mostly do not use the term *processing* at all, but that geo-processing definitions distinguish clearly between the two terms (e.g. ISO 19119). This also results in the fact that there is no consistent and uniformly applicable definition for both terms. Considering the goals and application scenarios of this thesis, the usage of the terms as follows can be defined as follows:

- *sensor fusion*: the enrichment of raw measurements towards "entities" so they can be used as input "datasets" for further processing steps
- *data fusion*: combination of several data sources with congeneric content (e.g. same data structures, objects, features etc.). This also includes basic data manipulation functionality (e.g. averaging, calculating mean values, simple filtering etc.).
- *geo-processing*: applying geospatial analysis algorithms on geo-referenced data to create meaningful and use case tailored information layers.

Within the Live Geography infrastructure – when operating on geographically referenced data – a differentiation between sensor fusion, data fusion and geo-processing can be defined, as shown in Figure 21.




Figure 21: Sensor Fusion and Data Fusion/Geo-processing in the Geo-Workflow Chain.

It shall be stated that besides real-time sensor measurements, the data fusion and the geo-processing steps will most likely also integrate legacy geospatial data like road networks, census data, or topographic models, and may also include historic/archived measurements. This is indicated by the databases ("DB") in Figure 21.

The output of the different processes can be summarised as follows. The result of a *sensor fusion* process is raw data, e.g. the O&M output of SOS data – harmonised data from several sensors of the same kind. The output of *data fusion* is data, e.g. the results of combining of several data sources, i.e. sensor fusion results, such as CO₂ and particulate matter concentrations combined from in-situ and EO sensors for different geographic areas. This also includes calibration and simple statistical analysis, i.e. mean value over time, variances, maximum and minimum values. On the contrary, *geo-processing* output can be seen as enriched information layers, i.e. analysed data. These can be air quality interpolation maps over a study area, where input data are sensor fusion results of CO₂, particulate matter and air temperature.

Figure 22 shows how sensor fusion, data fusion and geo-processing elements fit into the Live Geography workflow in terms of service chaining. The results of sensor fusion processes can be fed into data fusion and geo-processing operations, whereby different sensor fusion, data fusion and geo-processing operations can be

chained together, as shown in Figure 22. It shall be noted that the final output geographic feature set can again serve as an input for data fusion or geoprocessing operations, e.g. the harmonisation with the air quality index map of another city.




Figure 22: Chaining of Several Sensor Fusion and Data Fusion/Geo-processing Operations.

Connection to Complex Event Processing and Alerting

Complex Event Processing (CEP) is a method for detecting patterns in (sensor) data and thus serves for alert generation, for instance in case of threshold exceedance, quality control, or geo-fencing mechanisms, amongst others. Although CEP can also perform simple analysis operations by its inbuilt processing functionality, it is not considered part of the data fusion/geo-processing step.

CEP can be considered a possible interim step performing processing in the progression from raw measurements to data entities (cp. Figure 19). It can operate on both raw data, i.e. on untreated sensor measurements and on output of a sensor fusion process, and on fused/analysed data, i.e. output of data fusion or geo-processing operations, as long as its results are synchronously derived.

In the literature, CEP is sometimes considered a subset of sensor fusion. However, this conclusion is not valid for the Live Geography approach, as it would narrow down CEP's functional context to only being applicable within the sensor fusion step.

Within the Live Geography implementation presented in this thesis, CEP will be applied almost exclusively within and after the sensor fusion process (e.g. “send an alert if the pollutant concentration exceeds the threshold”). However, it can

theoretically also work on fused/analysed event data (e.g. "send an alert if the air quality indicator average exceeds the threshold over this particular area of interest"). Hence, CEP will be directly integrated into the geo-sensor network to allow event-based data processing in an early stage in the workflow chain.

CEP and related alerting methods will be described in more conceptual and technical detail in sub-chapter 6.6.

6.4 Geo-sensor Webs

Geo-sensor webs are the essential data source for monitoring applications. The design of a sensor network is defined by a number of parameters such as the operating environment, miniaturisation requirements, energy consumption, bandwidth utilisation or connectivity. These factors are highly application-specific and can thus not be uniquely specified for a general architecture like the Live Geography infrastructure.

Sub-chapter 5.1 illustrated a number of recent sensor network deployments, showing that research areas are highly heterogeneous; as well in sensor web research itself (technical issues as mentioned earlier in this paragraph) as in their application fields.

However, a fundamental prerequisite for urban monitoring networks is interoperability, which concerns data structures, measurement transmission, sensor queries and alerting functionality. Thus, the OGC Sensor Web Enablement (SWE) initiative has been chosen to assure interoperability in the Live Geography workflow. This choice was taken because of SWE's comprehensiveness (it serves the whole process chain), its broad development support as well from academia as from industry, its rapid advancement and introduction as official standards, and its maturity, which it has mainly gained through the last two years.

The most essential SWE standards for monitoring applications are the Sensor Observation Service (SOS) for sensor data provision, Observations and Measurements (O&M) for encapsulating measurements into a standardised XML-based format, and the Sensor Model Language (SensorML) for describing sensor platforms. Furthermore, the Sensor Alert Service (SAS) plays an important role for sending event-triggered alerts, e.g. in case of threshold transgression. These SWE standards have been described in more detail in sub-chapter 6.2.2.

The actual implementation of the geo-sensor web is highly application-specific according to the unique requirements of every single deployment. Thus, the Live Geography approach only suggests broad interoperability, which is achieved by the

extensive use of the OGC SWE standards, whereas other sensor network challenges have to be solved as the case arises.

6.5 Real-time Data Integration – Sensor Fusion

A substantial requirement for generating meaningful results from geospatial analysis processes is high quality of underlying data. As mentioned in Quality criteria for geospatial data comprise accuracy, completeness and topicality. While the first two parameters have recently received a lot of attention, the topicality parameter is still neglected in many cases.

This possibly arises from the fact that geospatial data is per definition historic and so far, geospatial analysis has focused on processing static data, with low temporal fluctuations. With the emergence of a variety of real-time data sources (e.g. sensor networks, RFID-based systems and Near Field Communication [NFC] systems etc.), and with the creation of the new paradigm of GIS-motivated situational awareness, topicality of base data is now rapidly gaining significance.

Thus, to enable the establishment of a live workflow as shown in Figure 10, also a mechanism for real-time integration of sensor measurements (sensor fusion) and other live data has to be created, which combines different data sources in real-time and provides them via standardised interfaces such as OGC Web Feature Service (WFS), Web Map Service (WMS), and Web Coverage Service (WCS). This allows for querying heterogeneous sensor data in real-time as well as their integration into GIS analysis systems.

As stated in sub-chapter 5.2, current real-time data integration approaches show various disadvantages. First, a substantial problem arises from the fact that data fusion originated in the domain of computer science. Hence, very few approaches exist, which are dedicatedly designed for geospatial data. Subsequently, integration of sensor data into GIS analysis systems at the moment mostly happens via the laborious interim step of a temporary physical database. This is not desirable in an automated GIS workflow as the database can easily become a bottleneck for large data volumes. Moreover, this approach unnecessarily adds another component to the overall workflow, which can result in substantial losses in performance. Thus, a need arises for an approach towards on-the-fly integration of sensor measurements.

Furthermore, current data integration systems are often inflexible for the adaptation to other data structures, while being characterised by a high degree of external dependence to particular data fusion software.

This is often the result of lacking compliance of the software with open data provision standards restricting response formats to very application-specific and proprietary data structures. This is a central issue particularly in integrating real-time sensor data into GIS products, which use a variety of well-established data structures, but as well outside the geospatial domain. Apart from that, output data in standardised formats can be easily integrated into mapping APIs such as Google Maps or Microsoft Virtual Earth, which have recently gained extensively in popularity.

The key requirement of the Live Geography approach towards sensor fusion presented within this thesis is data provision of real-time measurements for heterogeneous application domains and requirement profiles via well-established geo-service interfaces. That requires on-the-fly transformation of sensor data into standardised data provision formats, while enabling the establishment of a distributed service-oriented sensor web architecture.

The main advantage of this approach is that sensor data are made available through a variety of established geo-standardised interfaces (OGC's WFS, WMS, WCS), which offer a variety of output formats such as OGC Geographic Markup Language (GML), OGC Keyhole Markup Language (KML), geoRSS, geoJSON, Scalable Vector Graphics (SVG) or PDF.

The direct conversion of data structures allows for the easy integration of sensor data into GIS applications and therefore enables fast and ubiquitous data visualisation and analysis. This enables integration of output data not only in specialised GIS products, but as well outside the geospatial domain by using custom formats such as SVG or PDF. Apart from that, output data can be easily integrated into mapping APIs such as Google Maps or Microsoft Virtual Earth by using WFS output in geoRSS or geoJSON formats.

Naturally, there is a wide variety of viable implementations for this approach. One possibility is a custom data store for the open-source server GeoServer (OpenGeo, 2009), which uses O&M-encoded SOS measurements as input data and transforms them on-the-fly into various different output formats provided by OGC WFS, WMS and WCS service interfaces. More implementation-specific advantages of this sensor fusion approach are illustrated in sub-chapter 7.3.

6.6 Location-aware Event-based Alerting

Apart from standardised data transmission and provision, a special focus in the Live Geography approach is the extension of Complex Event Processing (CEP) functionality by spatial components. In a geographic context, CEP can e.g. serve for detecting threshold exceedances, for geo-fencing implementations, for investigating spatial clusters or as well for assuring data quality. To guarantee real-time measurement quality (e.g. to detect sensor failures or measurement outliers), location-enabled CEP technology can be used following Waldo Tobler's first law of geography stating that near features are more related than distant features. As CEP is not a well-known technology (as opposed to sensor networks and data integration), its basics are illustrated within this sub-chapter.

Generally speaking, CEP is a technology that extracts knowledge from distributed systems and transforms it into contextual knowledge. Since the information content of this contextual knowledge is of higher quality than in the original raw data, business decisions that are derived from it can be more accurate. A general architecture approach including five steps is shown in Job Description Language (JDL) in Figure 23, which has been adapted from Bass (2007).




Figure 23: General CEP Architecture with Data Processing Levels. (Adapted from Bass, 2007)

The CEP processing chain is divided into five levels that represent processing steps and the database management component, which is responsible for handling data and profiles. Since individual systems may have different requirements, the implementation depth will vary from system to system. The following list notes the basic requirements for each level, according to Bass (2007).

- **Level P: Pre-processing** – Before the actual processing, the input data is normalised, validated and eventually pre-filtered. Additionally, it may be important to apply feature extraction to eliminate data that is not needed for

the event filtering process. Although this is an essential step, it is technically not part of the CEP process and therefore marked as "Level P".

- *Level 1: Event Refinement* – This component's task is to track and trace an event in the system. An event's characteristics, such as data, behaviour and relationships, are translated into event attributes, i.e. prepared for further processing. The tracing part deals with state estimation that tries to predict upcoming events in the system.
- *Level 2: Situation Refinement* – The central part in the CEP workflow is the Situation Refinement, where the actual processing of simple and complex events takes place. This analysis includes mathematical algorithms, which comprise not only computing boundaries for certain values, but also matching them against patterns and historical data. The results are high level (contextual) interpretations, which can be acquired in real-time.
- *Level 3: Impact Assessment* – After having analysed and refined the situation, the next step is to generate decisions that may have consequences. Impact Assessment deals with the simulation of outcomes. It takes various scenarios and simulates them by taking cost factors and resources into account. The results are weighted decision reports that include priorities and proposals for accordant scenarios.
- *Level 4: Process Refinement* – Finally, the last level deals with the interaction between the CEP system and business processes. It defines a feedback loop to control and refine business processes. While this could include integration and automatic controlling of processes, e.g. within a geo-workflow, it may also just be the creation of alerting messages or business reports.

CEP has originated and traditionally been implemented in the financial and economic sectors to predict market developments and exchange rate trends. In these areas, CEP patterns emerge from relationships between the factors *time*, *cause* (dependency between events) and *aggregation* (significance of an event's activity towards other events). In a location-aware CEP these aspects are extended by additional parameters indicating spatial information of the event. This means that location parameters are combined on par with other relational indicators.

The new approach of integrating the "location" parameter in CEP systems, which is pursued in this thesis, allows for rule-based data processing of measurement data and for creating complex events accordingly. These can be related to time (temporal validity), space (e.g. geo-fencing with geographic "intersect", "overlap" or "join" operations) or measurement parameters (e.g. threshold exceedances).

Event recognition and processing happens in two different stages of the workflow. Firstly, at sensor level, CEP is used to detect errors in measurement values by applying different statistical operations such as standard deviations, spatial and temporal averaging or outlier detection. Secondly, after the data harmonisation process, CEP serves for spatio-temporal pattern recognition, anomaly detection, and alert generation in case of threshold violation.

As geo-referenced data is therefore subject for further processing, one important aspect is to control its data quality. Quality criteria comprise lineage, logical consistency, completeness, temporal quality, or spatial accuracy. From a practical viewpoint, this means that the location may be used to define quality indicators which can be integrated into CEP pattern rules. This happens by the use of external geographical code libraries, as presented in sub-chapter 7.4.

This integration of the location parameter allows for performing simple geographic "processing" operations directly within the CEP module. Within the Live Geography approach, it shall be used primarily for performing operations on events and synchronously derived event data (e.g. sensor fusion results). These operations comprise – amongst others – quality assurance like filtering and error detection, whereas more complex processing is performed by data fusion and geo-processing operations. In effect, this means that the CEP component can trigger a geo-processing/data-fusion operation based on an event, which appears in fused data.

Integration of SAS and CEP

As described in sub-chapter 6.2, the OGC Sensor Alert Service (SAS) specifies interfaces (not a service in the traditional sense) enabling sensors to advertise and publish alerts including according metadata. It shall be noted that "alerts" are not only understood in the classic meaning of the word, i.e. an automatic signal or notification indicating that an event has fired (e.g. a message in case of threshold exceedance), but in a broader context. Alerts are defined as "data" sent from the SAS to the client, which may as well comprise alerts/notifications (e.g. OGC Web Notification Service [WNS]) as observational data (measurements matching pre-defined criteria).

From an implementation viewpoint, SAS may be integrated with a CEP system as it operates asynchronously – after a synchronous initialisation process. An adapter for the service manages the subscriptions and converts received alerts into events that can be processed by the filtering engine. The adapter itself maintains a list of all established connections. A connection comprises several subscriptions and their corresponding chat rooms (MUC). In order to retrieve a sensor subscription, the

adapter sends a request. This request can be used for all three types of service operations – *subscribe*, *renew subscription* and *unsubscribe*. After the CEP process has finished, the request object fires a service event containing the result of the operation. If the operation was successful, the SAS adapter adds the subscription data to its list and joins the returned chat room.

Moreover, a data adapter is necessary, as only pre-processed information can be parsed by a CEP system, as shown in Figure 23. This adapter is responsible for data acquisition and transforming it into pre-defined event types. Moreover, a CEP module is necessary to connect the CEP engine with the data adapter and the resulting complex events. Within this thesis, a prototype implementation has been performed using the Esper CEP engine. This technology choice is motivated in sub-chapter 7.4.

For the integration of geographic topological functions in the filtering process, it is necessary to offer a possibility to define these functions when entering a CEP rule. The Esper CEP engine provides this functionality in the form of User-Defined Functions (UDF). A set of topological rules (based on open-source Java Topology Suite) has been implemented for OGC Simple Features, as described in sub-chapter 7.4.

Concluding, it shall be stated that SAS can be combined with CEP in a very modular fashion allowing for loose coupling of event processing and alerting/notification components. A future point of discussion is the integration of CEP and alerting with the OGC Web Processing Service (WPS). This issue is however not in the scope of this thesis.

Sensor Alert Service vs. Sensor Event Service

In contrast to SAS, the OGC Sensor Event Service (SES) standard specifies a service interface for filtering sensor data according to pre-defined filter criteria, whereby the filtering process itself operates on (sensor) data or on events. In this context, events are defined as "actions that occur at an instant or over an interval of time". SES defines three levels of filtering:

- syntactic filtering of sensor data
- filtering on single data sets (sensor measurements)
- filtering of data streams ("data fusion")

Another difference versus SAS is that SES does not define a transport protocol for performing filtering operations, whereas SAS specifies dedicated protocol usage; HTTP is used for advertising sensor data and client requests, whereas the

Extensible Messaging and Presence Protocol (XMPP) is used for push based alerting. More detail about the SES standard can be found in Echterhoff and Everding (2009).

Generally, SES (currently in version 0.0.1) is designed as a functional extension of SAS (currently in version 0.9) in terms of event (stream) processing. In addition to SAS, SES not only specifies the interface for advertising and publishing alerts, but also for filtering operations themselves. Consequently, this means that providing the same functionality like the SES, SAS has to be coupled with an event processing engine, which creates events. Again, after this processing, results can be published via the SAS interface over the XMPP protocol.

6.7 Real-time Geo-processing Using Server-based Analysis Technology

Besides geo-sensor networks and real-time sensor fusion, a major future challenge is the fusion (full integration) of real-time sensor measurement information, legacy GIS resources – data and services – and earth observation information in a "live" geo-aware environment. This requires an appropriate supply of input information streams from both space and terrestrial observation systems. It is necessary to identify effective ways for providing these data for fusion in a reliable way and with a guarantee of long-term continuity. The aim is also to support the future development of an appropriate European capacity in this context, coupled with a long-term perspective of sustainability as well as the integration feasibility into global system infrastructures like GEOSS (Global Earth Observation System of Systems).

The integration of real-time measurements into geographic modelling allows for introducing a reassessment and a new perspective of the quality criterion "topicality" in analysing and processing geographic phenomena in real-time – leveraged for environmental monitoring, safety, security, health care and a variety of other end applications.

The distinct advantages of server-based data analysis are threefold. Firstly, users do not need expert knowledge to use GIS tools as only simple web interfaces are provided to the end user; secondly, end users do not need to buy expensive and resource-consuming analysis software as processing requirements are shifted away from desktop computers to the server-side (naturally, this does not apply to not service providers); and thirdly, analysis models can be made accessible to a far broader spectrum of users by exposing them to the web, ranging from private

persons who want to know about air quality in their neighbourhood, to urban planners, traffic management officers, and political decision makers.

Summarising, the Live Geography approach shifts resource-consuming geo-processing operations away from the client by executing complex, asynchronous analysis tasks on the server side, and then simply providing the client with a tailored result. The output could be an XML structure, a numerical value, or a contextual map tailored to the user's specific needs. The crucial benefit of this approach is that GIS applications – that previously offered GIS functionality only through resource-consuming desktop clients – could be replaced by lightweight web-based analysis tools. This enables the results of geospatial analysis to be delivered to a broad spectrum of user-specific clients. It also allows for real-time situational awareness in spatial decision support systems. In other words, the infrastructure is suitable for integrating a wide variety of compliant data sets to perform use case specific geospatial analysis.

These benefits mentioned above will likely induce a paradigm shift in geospatial data processing in the next years and foster a broader spectrum of application areas for spatial analysis operations. Thus, the whole Live Geography approach is applicable to a wide variety of cross-domain use cases due to its high degree of flexibility and interoperability, as shown in sub-chapter 8.6.

Chapter Summary

Utilisation of urban real-time data in GIS applications requires a rethinking process of existing practices. Currently, geo-sensor networks are built up in closed single-purpose monolithic systems limiting interoperability and reusability of both data and workflow components. Moreover, geospatial data analysis is performed in costly specialised software involving a high degree of manual intervention for data gathering, pre-processing and quality assurance. Furthermore, geospatial analysis is per definition applied on historic data with long cycles from data generation to analysis output and a real-world impact.

The Live Geography approach seeks to tackle these challenges with an open sensing infrastructure for urban monitoring applications including web-based data analysis. The system makes extensive use of open (geospatial) standards throughout the entire process chain – from sensor data integration to analysis, Complex Event Processing (CEP), alerting, and finally information visualisation. It discusses the conceived modules as well as the overall created infrastructure as a whole.

As the Live Geography approach accounts for the entire workflow, it builds the architectural bridge between domain-independent sensor network development and use case specific requirements for end user sensitive information output.

The major Live Geography benefits are interoperability and component reusability through the broad use of open standards, the possibility for real-time data integration into GIS systems, and finally portability to a variety of application domains, which is achieved by the creation of a highly flexible infrastructure loosely coupling service-oriented building blocks.

IMPLEMENTATION

7 Implementation of the Live Geography Approach

The Live Geography approach is manifold in terms of functional components, as illustrated in chapter 6. Thus, there is a wide variety of implementation possibilities for each of these components (geo-sensor web, sensor data integration, and web-based data processing).

The implementation described within this chapter comprises tailor-made sensing devices, a real-time sensor fusion mechanism, a CEP and alerting module, a use case specific interpolation model and an automatic server-based analysis component. All these stand-alone modules are chained together using open standards as described below. The underlying scientific challenge is the conception of a framework, which allows for the integration of standardised pre-processed sensor measurements into GIS (Geographic Information Systems).

7.1 Embedded Sensing Device

The measurement device has been particularly designed for pervasive GIS applications using ubiquitous embedded sensing technologies. The system has been conceived in such a modular way that the base platform can be used within a variety of sensor web applications such as environmental monitoring, biometric parameter surveillance, critical infrastructure protection, or energy network observation by simply changing the interfaced sensors.

As it can be seen from Figure 24, the communication of the sensing device with other components in the workflow occurs on the basis of open standards of the Sensor Web Enablement (SWE) family. This requires a SensorML-conformal description of the measurement platform, O&M-compliant encapsulation of measurement values (as described in sub-chapter 7.2), as well as an SAS-compliant alerting module, which is separately explained in sub-chapter 7.4. In addition, a CEP engine implementing event-based data processing has been integrated directly on the embedded device.

Furthermore, an embedded database has been implemented directly on the sensor device to provide for the possibility of short-term data storage. This enables trend analysis and quality assurance, and reduces communication overhead with the central archive database.




Figure 24: Communication of the Sensor Device with other Components in the System.

The device also implements the following essential standards of the SWE family, as illustrated in Figure 24. These standards are described in more detail in sub-chapter 6.2.

- *Observations & Measurements (O&M)* – O&M allows for the formalised description of sensor measurements in a structured XML-based encoding schema. Thus, O&M can map sensor parameters and their relations. Measurements are organised by quantities, categories as well as their spatial and temporal characteristics. (Cox, 2007)
- *Sensor Model Language (SensorML)* – The Sensor Model Language (SensorML) is a general schema for describing functional models of the sensor. Information provided by SensorML includes observation and geometry characteristics as well as a description and a documentation of the sensor, and a history of the component's creation, modification, inspection or deployment. (Botts, 2007)
- *Sensor Observation Service (SOS)* – SOS allows for standardised access to sensor measurements (return type O&M) and their platform descriptions (return type SensorML) via a web service interface. (Na and Priest, 2007)
- *Sensor Alert Service (SAS)* – SAS is a service for the surveillance of pre-defined rules and trigger specified actions in a particular workflow in case of violation of these rules. (Simonis, 2007a)

The detailed software and service infrastructure running on the embedded sensing device is described in sub-chapter 7.1.2.

The pod connects sensors via a variety of hardware interfaces such as serial, UART (Universal Asynchronous Receiver/Transmitter), I²C (Inter-Integrated Circuit), Universal Serial Bus (USB), ZigBee, Bluetooth etc., stores a short history of measurement data in the embedded database and serves data over the internet using a secure web server (HTTP, HTTPS and XMPP). This standardised data provision can either happen using the SOS (providing measurement data in O&M format) or the SAS (providing alerts in an asynchronous fashion).

7.1.1 Sensor Pod Hardware

Compared to general-purpose computers (e.g. PCs), embedded devices are particularly characterised by their specific functionality and their real-time capabilities. Thus, embedded devices are employed in a wide variety of applications, ranging from MP3 players to industrial control devices, building automation, or automated diagnosis systems in the automotive sector. The main characteristics of embedded devices are thus:

- tailored design for a specific field of application
- integrative design
- minimised costs
- low energy consumption
- optimised software infrastructure
- constrained capabilities of the hardware infrastructure

The sensor pod used in this implementation consists of a COTS embedded device, a Gumstix⁴ Verdex XM4 platform including an ARM7-based 400MHz processor with 64MB RAM and 16MB flash memory. Generally speaking, Gumstix offers a highly modular and easily expandable system. The computer-on-module (the actual embedded device including CPU, memory and several interfaces) offers two I/O ports, which allows for extensibility of the basic system by specific modules such as GPS, Bluetooth, Wireless Fidelity (WiFi), Local Area Network (LAN), interface breakouts or a console board for programming the device.

⁴ <http://www.gumstix.com>

In the configuration for this thesis-specific implementation, a GPS module (U-BLOX NEO 4S and LEA-4P) has been attached for positioning and a breakout module for interfacing several different sensors (e.g. LM92 temperature, SHT15 for air temperature and humidity, NONIN 8000SM oxygen saturation and pulse, or SSM1 radiation sensors) via standardised interfaces like UART, I²C, USB etc. For the configuration of the device, a *console-vx* module (serial console connection) has been connected. This enables the establishment of a connection to the embedded device via the miniDIN8 RS232 port for installing the bootloader, the operating system and other software, as well as for debugging. In the operational process, this console is removed. The following figure shows the ready-to-deploy embedded device including the main computation component, a breakout module including the SHT15 air temperature and humidity sensor and a GPS module. A customary Universal Mobile Telecommunications System (UMTS) dongle for communication is attached via the integrated micro-USB port.




Figure 25: Gumstix Verdex Pro XM4 with GPS and a Combined Air Temperature and Humidity Sensor.

The size of the complete sensor pod is approximately 80x20x30mm, i.e. about the size of a chewing gum package. In full load, the device features an energy consumption of <2.2W including a running data query, the GPS module and data transmission via UMTS, which is known to be comparatively energy intensive way of broadcasting data. The following calculation yields an operation time under full load of 9.1 hours given a battery capacity of 4.000mAh.

$$Q \dots 4000mAh$$

$$P \dots 2.2W$$

$$U \dots 5V$$

$$P = U \cdot I \Leftrightarrow I = \frac{P}{U}$$

$$Q = I \cdot t \Leftrightarrow t = \frac{Q}{I} = \frac{Q \cdot U}{P}$$

$$\Rightarrow t = \frac{4000mAh \cdot 5V}{2.2W} = 9.1h$$

The electric charge is defined as the product of current conduction and time ($Q = I \cdot t$), which in turn means that the time “a battery lasts” is the electric charge divided by the current conduction ($t = \frac{Q}{I}$). As the current flow can be expressed as electric power over voltage ($I = \frac{P}{U}$), i.e. $1W=1VA$, time can be calculated by the term $t = \frac{Q \cdot U}{P}$.

The above calculation was done using an assumed electric charge of 4.000mAh, which is held by a reasonably-sized rechargeable Lithium-ion Polymer (LiPo) battery (140x40x10mm) – where its capacity and size is naturally dependent on the specific use case – and a constant voltage of 5V as per the specifications of the embedded device. In a production environment, the sensor pod shuts down non-active components (transmission unit, GPS receiver etc.) in order to increase battery lifetime.

7.1.2 Software Infrastructure

The sensing device runs a customised version of the “Open Embedded” Linux distribution (kernel version 2.6.21) with an overall footprint of about 8MB. This basic operating system houses a software infrastructure, which is comprised of several services, daemons and hardware drivers, as shown in Figure 26. The software infrastructure comprises an embedded secure web server (Nostromo nhttpd), an SQLite database and several daemons, which convert sensor readings before they are served to the web. The database acts as short-term storage of historic measurements to allow for different error detection procedures and plausibility checks, as well as for non-sophisticated trend analysis.




Figure 26: Software Infrastructure on the Embedded Sensing Device.

The hardware drivers for interfacing sensors and reading their measurements make up the first part of the embedded software infrastructure. As the geographical position is an essential must-parameter in geo-sensor networks, the sensor pod always interfaces a location sensor (e.g. a GPS/Galileo module, a ZigBee/WiFi-based positioning component etc.). Thus, *position* and *time*, which is necessary for a distinct identification of the temporal context, are two parameters that always have to be included into a measurement data set.

These measurements are then read by a special sensor daemon that essentially builds the bridge between the sensors and the internal software components. The daemon, which is always running in the background, interprets the bit stream coming from the sensor interface and converts it into meaningful measurement data.

These data are then stored into an embedded database (SQLite), which is held at a maximum data set volume, currently 12.500 readings. The database implementation and the relational data model will be described separately in sub-chapter 7.1.3. The sensor data, which is stored in the database, is then accessed from two different web servers (HTTP/HTTPS and XMPP), which make the measurements accessible from the internet.

HTTPS is considered a high enough security level for this implementation providing a secure channel between server and client using the Secure Socket Layer (SSL) protocol. Web Service Security (WSS), which has been described in sub-chapter 4.3, would be a viable alternative providing message-based security. However, as WSS is using the Simple Object Access Protocol (SOAP), it is characterised by large overhead, which is not well suitable for embedded sensor unit implementations.

The *Nostromo*⁵ web server was chosen as HTTP/HTTPS server because of its very compact design and its low resource requirements. Its minimal footprint of 50kB flash memory usage highly qualifies the server for use with embedded devices. Furthermore, Common Gateway Interface (CGI) functionality is integrated, which is essential for providing sensor data via the internet. A substantial advantage of Nostromo over the *boa* server, which comes pre-installed with the Gumstix operating system, is its support of secure connections via HTTPS without the necessity of installing additional plug-ins. Nostromo uses a 256 Bit Advanced Encryption Standard (AES) encryption for securing HTTP connections to a client.

Due to its small footprint and open-source availability, the *jabberd2*⁶ Extensible Messaging and Presence Protocol (XMPP) server was chosen for the implementation of alerting functionality on the embedded device. XMPP is an Internet Engineering Task Force (IETF) standard for push-based communications technology. XMPP was designed for real-time XML-based communication, enabling a broad spectrum of applications such as instant messaging, media negotiation, collaboration and more.

⁵ http://www.nazgul.ch/dev_nostromo.html

⁶ <http://jabberd2.xiaoka.com>

As alerting functionality including CEP and SAS implementations is a vital part of the workflow, it is treated separately in sub-chapter 7.4.

Finally, the “top” of the software stack in Figure 26 consists of two service interfaces towards the internet. The OGC SOS and SAS make sensor data available over the web. The actual embedded implementations of SOS and SAS will be described in sub-chapter 7.2.

7.1.3 Embedded Database

In order to allow for short-term storage of sensor measurements, the sensor pod runs an embedded database. It basically serves for organisation of a small amount of data sets (e.g. 12.500 as explained below) to enable quality control, error detection or trend analysis directly on the embedded device.

The requirements for the implementation of an embedded database are a minimal footprint and the support for the quasi-standardised Structured Query Language (SQL). Thus, SQLite was chosen, which is a “self-contained, serverless, zero-configuration, transactional SQL database engine” (Hipp, 2009). SQLite is an in-process database, meaning that it does not have a separate server process. Furthermore, the database file format is cross-platform qualified enabling deployment on various operating systems and hardware platforms.

The first realisation of an OGC-conformant Sensor Observation Service (SOS) and currently a reference implementation is done by the open-source initiative 52°North (Stasch, 2009). Figure 27 shows the used database model.




Figure 27: 52°North Relation Database Schema for Mobile Sensors. (Stasch, 2009)

As the database model in Figure 27 using 22 tables, lots of constraints and a tremendous degree of normalisation, it is by far too complex to achieve reasonable performance, particularly for implementation in embedded databases. Thus, it was

decided to create a separate schema for the Live Geography implementation, as illustrated in Figure 28.

The model is divided into three categories of sensor data sub-sets. Firstly, the *observation* part contains the observation itself and the feature of interest; secondly, the *measurements* section comprises the phenomena and procedures; and thirdly the *values* set covers the measurement values itself including their position. These terms are defined in sub-chapter 7.2.1.




Figure 28: Relational Schema of the Embedded Measurement Database Structure.

The embedded database schema has been designed towards simplifying SOS queries while aiming for an optimal trade-off between normalisation, completeness and performance. The element structure is highly similar to SOS process logic and to O&M data structures. For the Live Geography implementation, the database holds 12.500 measurement values, which approximately corresponds to 1MB of data volume presuming that one dataset is 80 bytes. This limit was set to account for the specific sensor device's limited storage capacities and can naturally be adapted for other use cases. The presented relational schema is reasonably scalable as it uses few constraints, but still corresponds to the 4th normal form. In order to prevent $m:n$ -relations, interim tables have been inserted to ensure data integrity.

To allow for longer-term data analysis, the measurements stored directly on the sensing device have to be transferred to an archive database, which is housed in the data centre as mentioned in chapter 6. However, this collection database has not been set-up within the implementation phase of this thesis as this is a standard process without major research aspect. Several spatial relational database products offer ready-to-use automated synchronisation mechanisms. Thus, it was not necessary to develop a dedicated software component for this purpose. However, it shall be noted that a database storing historical measurement data is naturally essential for sophisticated longer-term analysis in the context of fully-fledged Spatial Data Infrastructures (SDI).

7.2 Standardised Service Infrastructure

For standardised data retrieval, a service has been created implementing the OGC Sensor Observation Service (SOS), Sensor Model Language (SensorML) and Observations and Measurements (O&M) standards in an application-specific way. Like this, measurement data are served via HTTP/HTTPS over UMTS in the standardised XML-based O&M format over the SOS service interface. SensorML is used to describe the whole sensor platform as well as to control the sensor via the OGC Sensor Planning Service (SPS), e.g. to dynamically adjust measurement cycles. It shall be mentioned that SPS has not been implemented within this thesis. Finally, SAS is responsible for creating events and alerts from measurement data. The overall service stack on the embedded sensing device is illustrated in Figure 26.

7.2.1 Embedded OGC Sensor Observation Service

The OGC Sensor Observation Service (SOS) serves for querying measurement data from the embedded sensing unit. The service implements the three mandatory methods as specified by the OGC SOS standard (Na and Priest, 2007), *DescribeSensor*, *GetCapabilities* and *GetObservation*. Basically, the service, which is implemented in CGI, parses the request and creates the according response using appropriate XML templates.

The *DescribeSensor* operation returns a SensorML-conformant description of the sensor platform. The *GetCapabilities* method provides a SOS capabilities document including identification, input and output lists, contact data and measurement metadata, amongst others. The *GetObservation* function queries the embedded database according to the *GetObservation* request and returns the required measurement data in the standardised O&M format. The requests and responses for

the described operations are listed in the Annex, as these documents are fairly comprehensive. The general O&M data structure is shown in Figure 29. It shall be mentioned that this is a simplified view on the connection of O&M elements to clarify their basic interplay, without claiming completeness as specified by the OGC O&M standard.




Figure 29: OGC Data Model for O&M.

According to the OGC O&M standard (Cox, 2007), the elements in Figure 29 are characterised as follows:

- *observation* – an act of observing a property or phenomenon, with the goal of producing an estimate of the value of the property. A specialised event whose result is a data value.
- *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.
- *property* – characteristic of a feature type, including attribute, association role, defined behaviour, feature association, specialisation, and generalisation relationship and constraints.
- *procedure* – method, algorithm or instrument. It is the description of a process used to generate the result. It must be suitable for the observed property.

- *result* – an estimate of the value of some property generated by a known procedure.
- *featureOfInterest* – a feature (abstraction of a real-world phenomenon) of any type, which is a representation of the observation target, being the real-world object, regarding which the observation is made.

7.2.2 Embedded OGC Sensor Alert Service

The implemented Sensor Alert Service (SAS) is responsible for sending alerts, which have been triggered by events created in the CEP engine, as explained in sub-chapter 7.4. The service implements the mandatory operations as specified in the standard, namely *DescribeSensor*, *DescribeAlert*, *GetCapabilities*, *Subscribe*, *RenewSubscription* and *CancelSubscription*, as illustrated by Simonis (2007a).

The *DescribeSensor* and *GetCapabilities* operations are used analogically to the SOS methods described in sub-chapter 7.2.1. The *DescribeAlert* method allows users to request a template of the alert message structure.

The *Subscribe* operation allows users to subscribe to a particular sensor alert, i.e. to receive a notification if a pre-defined rule applies. If a user registers for an alert, the sensor returns a *SubscribeResponse* indicating the XMPP channels used for messages and acknowledgments. Moreover, the response includes an expiration timestamp, until the data may be retrieved, and a unique ID for identifying the subscription.

As its name indicates, the *RenewSubscription* operation allows clients to renew their subscription to a sensor alert before it expires. In case such a request is issued, the sensor automatically extends the expiration time and responds with a *RenewSubscriptionResponse*. Accordingly, a client can also cancel a subscription by sending a *CancelSubscription* request.

A more detailed description of the SAS implementation can be found in sub-chapter 7.4, as this module is closely coupled with the Complex Event Processing component. Again, the request and response documents for the described operations can be found in the Annex.

7.3 Real-time Data Integration – Sensor Fusion

As mentioned in sub-chapter 6.5, the significance of topicality in base data for geospatial analysis processes is rapidly growing with the emergence of a variety of real-time data sources.

Thus, a mechanism for real-time integration of sensor measurements (sensor fusion) and other live data has to be created, which combines different data sources on-the-fly and provides them via standardised interfaces such as OGC Web Feature Service (WFS), Web Map Service (WMS), and Web Coverage Service (WCS). This allows for querying heterogeneous sensor data in real-time and facilitates their integration into geospatial analysis systems.

7.3.1 Technical Realisation

For the sensor fusion component, a data store extension to the open-source product *GeoServer* (OpenGeo, 2009) 1.6, has been developed using Java technology. This plug-in enables the direct integration of OGC SOS responses (in O&M [Observations and Measurements] format) into *GeoServer* and their conversion to OGC-conformal service messages on-the-fly. The main advantage of this approach is that sensor data are made available through a variety of established standardised geo-service interfaces (OGC standards WFS 1.0, WMS 1.1.1 und WCS 1.0), which offer a variety of output formats such as OGC Geographic Markup Language (GML), OGC Keyhole Markup Language (KML), geoRSS, geoJSON, Scalable Vector Graphics (SVG) or PDF.

7.3.2 Functionality

During the transformation procedure from O&M input to WFS output, certain input parameters (coordinate reference system, unit conversions, data structures etc.) are interpreted. In practice, this means that the O&M XML structure is converted into well-established standardised data formats as mentioned above.

As described above, the O&M data store for *GeoServer* constitutes the connecting component between the geo-sensor web (Sensor Observation Service) and the standardised WFS interface, which is visible toward the internet for potential clients. Figure 30 illustrates the data flow sequence within the data store for a feature query in the form of method calls between the different classes of the data store.

The classes *SOSHashMap* and *MapCreator* are responsible for the actual data query, i.e. the communication with the external geo-sensor network, whereas the other

classes serve for integrating sensor data into GeoServer and for transforming O&M-encoded measurement data into WFS-conformant output formats.

The sequence diagram shown in Figure 30 is divided in five segments, as described in Table 2.

Segment	Description
1	GeoServer requests the "schema" of the queried feature type. Like this, GeoServer gains knowledge about the data structure – meaning the contained elements and their data types – and can thus dynamically build an accordant data query.
2	GeoServer creates a <i>Datastore</i> object, which is later used for the actual data query.
3	The data store calls the <i>getMap()</i> method, which queries the sensors in a parallel manner by making use of multi-threading technology. The measurement values are then assembled in a combined <i>HashMap</i> , which is returned to the data store.
4	For the actual operation of processing sensor data, the data store generates a <i>FeatureReader</i> object, which again creates an <i>AttributeReader</i> object. The latter reads the measurement data sets one by one, disassembles the single attributes according to the data schema, which has been requested in step 1, and saves them in the <i>FeatureReader</i> object.
5	The data store creates a <i>FeatureCollection</i> from the attribute data obtained in step 4 and returns it to GeoServer. From this data structure, the server is able to produce a variety of different output formats, such as WMS images, geoRSS, KML, coverages, PDF and others.

Table 2: Segments in the Sequence Diagram for the GeoServer Data Store Sensor Query.




Figure 30: Sequence Diagram of a Feature Query within the Custom Data Store.

7.3.3 Innovation and Benefits

The innovation of the GeoServer custom data store in comparison to previous data integration approaches is that the conversion of the data structure from SOS responses to various WFS, WMS and WCS output formats is performed on-the-fly.

This solution has a distinct advantage versus other recent approaches. First, conventional data integration methods typically use the laborious interim step of storing data in a temporary physical database. This is not preferable in an automated real-time workflow as the database can easily become a bottleneck for large data volumes. Second, this approach adds another component to the overall workflow, which can result in substantial losings in performance. This creates a single central point of failure making the system vulnerable in case of technical malfunctions. Thus, an approach towards on-the-fly integration of sensor measurements was chosen within this thesis.

Consequently, the Live Geography approach allows for the establishment of a distributed sensor service architecture, and thus enables data provision of real-time measurements for heterogeneous application domains and requirement profiles. The direct conversion of data structures allows for the easy integration of sensor data into GIS applications and therefore enables fast and ubiquitous data visualisation and analysis.

A substantial advantage of the GeoServer plug-in is manifested in its high flexibility and its low external dependence due to open access to GeoServer source code. Furthermore, the compliance of the server with three OGC data provision standards (WFS, WMS and WCS) offers a broad spectrum of different response formats. In effect, this enables direct integration of measurement data not only in specialised GIS products, but as well outside the geospatial domain using custom formats such as SVG or PDF. Furthermore, output data provided via standardised geo-service interfaces can be easily incorporated into mapping APIs such as Google Maps or Microsoft Virtual Earth by using in geoRSS or geoJSON output formats.

Another benefit of this approach is that no use case motivated and thus functionally limited web portal has to be implemented for sensor data access. In fact, live measurements can be queried via well-established OGC interfaces.

7.4 CEP and Location-aware Complex Event Processing and Alerting

Furthermore, a CEP and alerting mechanism based on Extensible Messaging and Presence Protocol (XMPP), conformant to the OGC Sensor Alerting Service (SAS) specification has been implemented.

For generating alerts, the OGC Sensor Alert Service (SAS) standard (s. sub-chapter 6.2) has been implemented for mobile sensor devices. SAS, which is part of the SWE initiative, specifies interfaces allowing sensors to advertise and publish alerts including according metadata.

Alerts, which are based on SAS Filter rules, are defined as “data” sent from the SAS to the client or to a Complex Event Processing Engine (CEP). These data can either be alerts/notifications (e.g. OGC Web notification service [WNS]) or sensor data (measurements matching pre-defined criteria). As SAS is based on the standardised XMPP protocol, alerts can be broadcasted very efficiently over the internet to subscribed consumers. (Simonis, 2007a)

Within the Live Geography concept, CEP is used for detecting patterns in measurement data and for creating complex events accordingly. These can be related to time (temporal validity), space (e.g. geo-fencing with geographic “intersect”, “overlap” or “join” operations) or measurement parameters (e.g. threshold exceedances).

These CEP operations can occur at two different stages of the workflow. Firstly, at sensor level CEP is used to detect errors in measurement values by performing different statistical calculations including standard deviations, spatial and temporal averaging, or outlier detection. Secondly, after the data harmonisation process CEP serves for spatio-temporal pattern recognition, anomaly detection, and alert generation in case of threshold transgression.

7.4.1 CEP Engines

As the CEP engine is the central component of the event processing module and only needed to be extended by geographical operations, an existing solution has been used within this implementation. Available CEP engines range from proprietary software, such as Oracle CEP (Oracle Corporation, 2009), to open-source projects, like Intelligent Event Processor (IEP) (CollabNet Inc., 2009b) or Esper (EsperTech Inc., 2009). The three mentioned engines have been evaluated according to pre-defined criteria, as shown in Table 3.

Regarding the license type, the IEP and Esper engines offer wider usage possibilities as they are delivered by open-source licenses that allow full usage for

research purposes. Apart from licensing issues, an essential factor is in what way the CEP frameworks can be integrated into an application. All three evaluated engines are accessed by a Java Application Programming Interface (API). Although basic access is identical, they use different ways of parsing event patterns. This means that each framework uses a non-standardised SQL-like syntax for characterising pattern rules. Oracle and Esper have defined proprietary specifications, whereas IEP employs the Continuous Query Language (CQL), which has been created at Stanford University (Arasu et al., 2003). The intended use of the CEP engine on a mobile device entails that the CEP toolkit is packaged in a compact way, e.g. as a single Java library. Currently only Esper can be embedded in this way into an application.

	Oracle CEP/EDA	IEP	Esper
Version	10.3	-	3.0.0
Type	Proprietary	Open-Source	Open-Source
License	-	CDDL	GPL v2
Language	Java (≥ 1.5)	Java (≥ 1.5)	Java (≥ 1.5)
EPL	Oracle Event Processing Language, graphic interface	CQL, graphic interface	Esper EQL

Table 3: CEP Frameworks.

7.4.2 Geographical Code Libraries

Geographical operations can be embedded into the geo-filtering process by importing an external library. Preferably, this should happen via the simple integration of a single Java based geo-framework. Currently, two functionally comprehensive software libraries are available: GeoTools (Open Source Geospatial Foundation, 2009) and the Java Topology Suite (JTS) (Vivid Solutions, 2009). They are comparatively evaluated in Table 4.

Both libraries are released under the GNU Lesser General Public License (LGPL), which makes them freely usable for research and in commercial products. The main differences between the two frameworks are their complexity and scope. JTS mainly focuses on topological functionality and has been created conforming to the OGC Simple Features Specification (SFS) for Structured Query Language (SQL)

(Herring, 2006). On the contrary, GeoTools is developed in compliance with the well-established OGC standards for data provision (WFS, WMS and WCS) to allow for more complex geographic operations on a set of information layers. In fact, GeoTools builds upon JTS for operations processing geometries and topologies. Additionally, GeoTools provides features, such as grids and data storage functions. Due to GeoTools' functional complexity, it requires several libraries, which supply the needed functionality, whereas JTS is packaged into a single library.

	GeoTools	JTS
Version	2.5.4	1.8.0
Type	Open-source	Open-Source
License	LGPL	LGPL
Language	Java (≥ 1.5)	Java (≥ 1.5)
Conformance	OGC	SFS for SQL

Table 4: Geographical Code Libraries.

7.4.3 Technical Realisation

For sending events created by the CEP engine, a push-based OGC SAS compliant alerting service has been realised. In this case, SAS is an asynchronous service connecting a sensor in a network to an observation client. In order to receive alerts, a client subscribes to the SAS. If specified rules apply, a pre-defined alert is sent to the client via XMPP. It shall be stated that the whole communication between the embedded XMPP server (jabberd2) and the client is XML-based for simplifying M2M messaging.

In the actual implementation, the Esper CEP engine in its version 3.0.0 has been used because of its open availability, Event Query Language (EQL) based event description, and its simple integration by just including a single Java library. The evaluation of available geographic code libraries yielded that JTS is the preferable solution for this implementation because of its easy integration. This choice has been taken despite GeoTools offers a richer feature set. However, JTS' features are more than enough for location-aware CEP implementations.

Figure 31 shows the internal sub-parts of the CEP-based event processing component, which is built up in a modular structure.




Figure 31: Internal Service Stack of the CEP Component.

Generally speaking, the event processing component connects the data layer (i.e. sensor measurements), and the data analysis and data visualisation components. In other words, it prepares raw data in order to be process-able in the analysis and the visualisation layers.

The *data transportation* component is responsible for connecting the data integration layer to a wide variety of data sources, which can comprise sensor data, real-time Really Simple Syndication (RSS) feeds, File Transfer Protocol (FTP) services, web services, databases etc. Hence, it serves as an entry point into the system. Its main responsibility is to receive various kinds of data structures and pass them on as is to the data transportation layer adding metadata to the actual payload, such as the data source or its format.

The *data transformation* component receives the data stream from the transportation module and parses it. Therefore, it has to implement a mechanism to treat the data format in the right way. This component parses the data in a pre-defined way and transforms it into objects, i.e. a commonly understandable data format or a class instance. After this conversion procedure, it passes the object on to the event processing engine or the data persistence component depending on pre-defined criteria such as the type of object, its origin or the data format.

The *event processing* component handles the objects coming from the transformation module according to specified user query statements. This module basically manages multiple streams, and identifies and selects meaningful events.

Currently, the Esper event processing engine is used to detect events and push data to the user processing component. The Esper engine has been extended by spatial parameters as described above. Following the event processing step, a further *user-defined processing* method can optionally be applied. This component implements a set of filtering or selection rules, which are specified by the user according to a very particular need.

The *data persistence* component receives a set of data from the data transformation module or from the processing modules (i.e. a “filtered” dataset). From these data, it creates a physical data structure, which can either be temporary (for rapidly changing real-time data) or permanent (for time-insensitive data, e.g. for stationary sensor locations). Consequently, as well live data sources as static ones can be handled by the data persistence component.

The *non-standard service* is one of three service interfaces, which connect the data integration layer to the data analysis layer. It provides data via a custom (i.e. non-standard) interface. This is necessary as existing standardised OGC services do not automatically support push mechanisms. Naturally, standardised data structures can also be served via the non-standard interface, just the service interface itself is not wholly standards-conformal. The service is also responsible for converting objects, which are created by the data transformation component, to a common, pre-defined output format. The output can either be e.g. GML, KML or geoRSS for spatial data, or RSS, JSON, SOAP bindings or a custom API for non-spatial data.

The *OGC service* is the second component, which connects the data integration layer to the data analysis layer. It offers a well-known and widely spread standardised interface in proved data structures such as GML, geoTIFF, KML or SVG. In essence, the main difference compared to the non-standard service is that this component provides OGC-standardised data structure over standardised service interface instead of providing a custom interface. The indicated *OGC real-time adapter* is basically a technological bridge to integrate live data into existing OGC services as existing implementations only support a variety of static data sources such as shape files, ASCII grids, different (geospatial) databases or cascading OGC services (WFS, WMS, WCS). Its implementation has been described in sub-chapter 7.3.

The *web service* component is the third interface connecting the integration layer with the analysis layer. It is intended for handling non-geographic non-real-time data and for serving it via the http protocol. This component is basically a regular web service, meaning that it implements the request/response communication model, but no pushing mechanism as the non-standard service component does.

7.5 Live Geospatial Data Analysis

The aim of the analysis component is the proof and verification of the integration of real-time measurements into GIS processing tools. This maximises significance for the different interest groups by integrating “live” information layers into complex geospatial processing and interpolation algorithms such as Inverse Distance Weighting (IDW), kriging or co-kriging to reveal unseen (spatial) correlations between different parameters.

Technically speaking, a framework for ESRI ArcGIS has been developed to integrate location-aware sensors directly into the GIS software product. To organise and reduce the vast amount of measurement information, OGC SOS services provide filtering capabilities in space, time and entity when leveraging this data source.

The user interface design has been limited to simple customisations. In a real-world deployment, this process obviously has to be performed specifically for each use case, including sophisticated methods of Human Computer Interaction (HCI) design. Chapter 8 includes the description of a number of end applications implemented within this thesis.

7.5.1 Desktop GIS Processing of Sensor Information

Since reliable processing services performing geospatial analysis tasks, which are based on open standards such as the OGC Web Processing Service (WPS), are not yet available, geospatial analysis is currently mostly performed in legacy desktop tools. Thus, a *Live Sensor Extension* has been created within this thesis, which allows for the direct integration of standardised sensor data into ArcGIS. The *Live View* component enables ad-hoc geospatial processing and visualisation in a navigable map using Dynamic Display technology.

Figure 32 shows an interpolated temperature surface in ESRI ArcMap using the IDW algorithm. In addition to ArcGIS, other free clients such as Open Layers, uDig, Google Maps, Google Earth or Microsoft Virtual Earth have been used to visualise sensor information in 2D and 3D.




Figure 32: 2D Interpolation Using the Inverse Distance Weighting (IDW) Algorithm.

7.5.2 Time-series Analysis

Apart from simple data visualisation, this plug-in allows for using sensor data in spatial analysis operations, which are offered by the ArcToolbox. Hence, all processing functionality, which is offered by ArcGIS Server, is available within the plug-in. Furthermore, complex and customised analysis models can be integrated by the use of ESRI Model Builder software.

Besides these spatial analysis tools, the Tracking Analyst extension developed within this thesis allows for time-series visualisation and temporal pattern recognition in measurement data, as illustrated in Figure 3.




Figure 33: Spatio-temporal CO₂ Data Analysis Using Tracking Analyst.

As mentioned above, this processing component uses ArcGIS Tracking Analyst tool to perform a spatio-temporal analysis on measurement data over a period of time. In order to achieve a coarse overview of pollutant variability, CO₂ data captured by the CitySense network in Cambridge MA, USA has been used. This allows for correlating temporal measurement data fluctuation to traffic density, weather conditions or day-time related differences in a very flexible way. The detailed functionality of this application is described in sub-chapter 8.6.

7.5.3 Web-based Geospatial Processing

A particularly innovative part of the implementation is the web-based geospatial processing component. Two data analysis models for Inverse Distance Weighting (IDW) and Kriging operations have been realised. Together with the web interface source code, these models were then integrated into a single toolbox, which can be published as a web service on ArcGIS server. This allows for data analysis by just selecting base data and the according processing method in the web client.

Two distinct advantages of this mechanism versus previous desktop GIS solutions are that geographic analysis can be done without profound expert knowledge and that processing requirements are shifted away from desktop computers to the server-side. These benefits will likely induce a paradigm shift in geospatial data processing in the next years and foster a broader spectrum of application areas for spatial analysis operations, as discussed in chapter 9.

Figure 34 illustrates with a semi-transparent overlay of the web-based IDW interpolation results of (air temperature data) with an orthorectified background image (accessed via OGC WMS) in a custom web client.




Figure 34: Web-based Interpolation of Real-time Sensor Data.

Conclusion

Currently, geospatial processing is mostly done in expensive and resource-consuming desktop applications. This is partly due to the fact that reliable and functional (open-source) server-based processing software is yet not widely implemented and still in development, which severely limits the choice of appropriate technology. Thus, the implementation for this thesis focused on ArcGIS Server software, which is particularly important for urban use cases, as public and governmental institutions often employ ArcGIS for geospatial analysis.

This real-time data integration approach using standardised Sensor Observation Services (OGC SOS) can be able to solve several important issues in the area of urban monitoring. The deployment of an open data integration and processing infrastructure is not only the prerequisite for a paradigm shift in collaboration and cooperation efforts between different user groups, but it is also likely to raise awareness for ubiquitous multi-domain sensing systems.

Chapter Summary

The Live Geography approach is manifold in terms of functional components. The implementation described within this chapter comprises tailor-made sensing devices, a real-time sensor fusion mechanism, a use case specific interpolation model and an automatic server-based analysis component. All these stand-alone modules are chained together using open standards. The underlying scientific challenge is the conception of a framework, which allows for the integration of standardised pre-processed sensor measurements into Geographic Information Systems (GIS).

In its current implementation, the Live Geography approach builds the basis for longer-term developments towards pervasive urban monitoring networks. However, the implemented components have to be taken further towards applicability in a production environment, which is a matter of source code consolidation and extensive testing. Furthermore, the infrastructure itself has to be extended by additional components to create a fully automated service infrastructure.

RESULTS

KE20FL2

8 Results

The main aim of this thesis is the development of a standardised measurement and analysis infrastructure approach for monitoring applications using distributed sensor networks with a particular focus on urban environments. A key requirement is maximum flexibility and modularity of the architectural design, so that several possible real-time geo-analysis applications can be implemented upon it. Furthermore, a sample application is implemented for demonstration purposes. Apart from these technical outcomes, possible impacts on urban dynamics and processes are briefly discussed.

This chapter links back to the main question, and to the strategic and operative research questions, by answering them according to the outcomes of this thesis. Chapter 2 defines the principal research question as:

How can interoperable infrastructures, which flexibly integrate geo-sensor networks, real-time sensor fusion and geospatial data processing, influence city dynamics by fostering the deployment of monitoring applications in urban environments and by complementing or even replacing monolithic measurement systems?

This overarching question can be sub-divided into several strategic questions and research task oriented operational questions:

Strategic Questions

- Why is interoperability an essential requirement for monitoring architectures and how can it be accomplished for geo-sensor networks?
- Which modules does a flexible geo-sensor data analysis workflow have to contain?
- How can the functionality of the infrastructure be expressed in terms of a workflow chaining together different service interfaces using (geospatial) standards?
- How can real-time data integration be achieved for geo-workflows?
- Which implications does “pervasive urban monitoring” contain for the individual?

- What societal, technological and perceptual impacts do pervasive monitoring infrastructures have on the urban environment?

Operational Questions

- What are monolithic measurement infrastructures and how do they limit sensor network development and deployment?
- What does the term “city” mean in terms of a functional application context for monitoring applications?
- What are particular challenges of urban monitoring?
- Which geospatial standards have to be taken into account for the creation of an interoperable measurement and data analysis infrastructure?
- What are the benefits of open interoperable sensing architectures?
- Which effects on the city and its inhabitants can potentially be expected using such systems?

8.1 Results Summary

This thesis presents the *Live Geography* approach proposing a technical infrastructure enabling a wide variety of urban monitoring applications. It suggests the establishment of an **interoperable, modular and flexible distributed sensing and data analysis infrastructure** particularly targeted to urban environments – as opposed to previous monolithic “stovepipe”-like sensor networks.

The dissertation comprises an illustration of technological background topics (standardisation, disambiguations of geo-workflow components etc.), as well as a description of the key requirements of the architectural modules. These modules comprise geo-sensor webs, real-time data integration and sensor fusion, location-aware event processing and alerting, and web-based geo-processing.

To prove the system’s portability and flexibility, **instantiations in different application domains** (live pollutant monitoring for public health, fine-grained air temperature variation assessment, ubiquitous biometric parameter surveillance and real-time air quality assessment) have been implemented using the same underlying infrastructure, as described in sub-chapter 8.6.

Furthermore, this thesis discusses how ubiquitous monitoring systems can **influence the city** – in its functional context comprising dynamic processes – in

real-time by feeding back analysed sensor information into the urban system (s. chapter 9).

As the Live Geography concept has been developed and implemented together with public partners as concrete end users, it is likely that the approach can **respond to dedicated needs of the city management**. Therefore, the longer-term goal is to augment people's perception of their environment by adding previously unseen information layers and thus influencing their short-term behaviour by providing real-time decision support.

8.2 The Urban Environment as the Functional Context

The significance of the definitions of the term "city" (s. sub-chapter 3.2) for this thesis is that the city as a functional construct shall not only be seen as a place of social interactions, financial transactions, a network of technology nodes, a geographical agglomeration area or as a political landscape, but more as an actuated multi-dimensional conglomerate of heterogeneous processes, in which the citizens are the central component. In other words, the **city can be seen as a complex near real-time control system** (Resch et al., 2008), creating a feedback loop between the city itself, the city management and the citizens, which is achieved by ubiquitous sensing of a variety of parameters. Sub-chapter 3.2 discusses the city as the functional context for this thesis in detail.

The urban context poses several specific **challenges** to pervasive monitoring: apart from well-known general monitoring challenges, very particular issues arise for geo-sensor network deployments in the city. These range from physical sensor mounting to social and privacy implications. Furthermore, the sensitive urban political landscape has to be accounted for, which might cause unforeseen challenges.

From a more contextual standpoint, an important peculiarity of the urban context is that there are large variations within continuous physical phenomena over small spatial and temporal scales.

Another central issue in deploying sensor networks in the city is the **impact of fine-grained urban monitoring**, as terms like "air quality" or "pollutant dispersion" are only surrogates for a much wider and more direct influence on people, such as life expectation, respiratory diseases or quality of life. This raises the challenge to find the balance between providing pervasive real-time information while still preserving people's privacy. As a result, it is self-evident that the provided information has to be highly accurate, reliable and unambiguous.

8.3 Live Geography Approach

Current implementations of urban geo-sensor networks are mostly built up in a monolithic manner in **closed single-purpose systems** due to heterogeneous interests of deploying and operating institutions. Thus, leveraging sensor data in an ad-hoc fashion is non-trivial as measurement data are mostly not combinable because of lacking interoperability and deficient coordination, as described in more detail in sub-chapter 1.3.

The Live Geography service infrastructure as illustrated in Figure 35 tries to tackle these challenges with a **fully standardised service-oriented module chain**, whereby standardisation is achieved by making use of well-established OGC standards such as Sensor Web Enablement (SWE), Web Feature Service (WFS), Web Map Service (WMS) or Web Coverage Service (WCS).

SWE is a new initiative aiming to make sensors discoverable, accessible and controllable over the internet. It currently consists of seven standards and interoperability reports, namely the Sensor Observation Service (SOS), Sensor Model Language (SensorML), Observations and Measurements (O&M), Transducer Markup Language (TML), Sensor Alert Service (SAS), Sensor Planning Service (SPS) and the Web Notification Service (WNS).

The general workflow within the infrastructure can be followed from left to right in Figure 35. First, heterogeneous data sources, such as sensor data, external data (provided via standardised interfaces such as OGC WFS or WCS), or archived data are integrated on the server side in a flexible sensor fusion process. This integration can happen via both classical request/response models and push services that send out alerts e.g. in case a certain threshold is exceeded. Next, the different kinds of data are combined by a data integration server. This step requires real-time processing capabilities, such as Complex Event Processing (CEP) or Event Stream Processing (ESP). The harmonised data are then fed to pre-defined geospatial process models to generate user-specific output.




Figure 35: Live Geography: Standards-based Loosely-coupled Service Infrastructure.

As the Live Geography approach covers the entire workflow from geo-sensor networks to user-specific outputs, it builds the **architectural bridge** between domain-independent sensor network development and use case tailored requirements for end user sensitive information output.

The major benefits of the approach are **interoperability and component reusability** through the broad use of open standards, the possibility for **real-time data integration** into GIS systems, the support for **server-based processing** technologies, and finally its **portability** to a variety of application domains, which is achieved by the creation of a highly flexible infrastructure loosely coupling service-oriented building blocks. These benefits are elucidated in more detail in sub-chapters 8.4 and 8.6.

It shall be stated that no profound security framework for the Live Geography has been created as there are already numerous solutions available, which have to be implemented as the case arises. The specific implementation presented within this thesis relies on HTTPS (Hypertext Transfer Protocol Secure) for secure data transmission from the sensor network to the sensor fusion server, as described in sub-chapter 7.1.

8.4 Benefits of the Developed Approach

Since the Live Geography system architecture presented in this thesis is composed of **layers of loosely-coupled and service-oriented building blocks**, sensor fusion can be decoupled from CEP, data analysis and visualisation components, allowing for flexible and dynamic service chaining. In order to fulfil real-time data processing needs and alerting requirements, the concept also incorporates an event-based push mechanism.

As noted above, one of the major challenges is the integration of location-enabled real-time measurement data into GIS service environments to perform distributed analysis tasks. There are three main requirements for quality-aware geospatial analysis: **accuracy, completeness and topicality** of the input parameters (layers). Previous developments often do not account for time stamp parameters. The Live Geography approach identifies effective ways to combine terrestrial and space real-time observation data with SDI information layers.

This combination of heterogeneous data sources for generating highly meaningful analysis results also involves the integration of legacy sensor networks and geospatial data. The Live Geography approach solves this issue by the wide use of **open standards**. This allows for providing existing real-time data via well-known service interfaces (OGC SOS, WFS, WMS or WCS) by implementing “proxy”

services, which encapsulate underlying data into pre-defined GIS formats. This enables simple integration of virtually any geospatial data source into analysis and decision support systems.

Another central benefit is that the Live Geography approach shifts resource-consuming geo-processing operations away from the client by executing complex, asynchronous analysis tasks on the server side, and then simply providing the client with a tailored result. The output could be an XML structure, a numerical value, or a contextual map tailored to the user's specific needs.

Thus, a crucial advantage of this approach is that GIS applications – that previously offered GIS functionality only through resource-consuming desktop clients – could be replaced by **lightweight web-based analysis tools**. This enables the results of geospatial analysis to be delivered to a broad spectrum of use case tailored clients.

Subsequently, this enables real-time situational awareness in spatial decision support systems. In other words, the developed infrastructure is suitable for integrating a wide variety of geo-data sets to perform use case specific geospatial analysis in real-time. This provides users with an up-to-date Common Operation Picture (COP) to base their decisions on.

Summarising, the major benefits of server-based data analysis are that (a) users do not need expert knowledge to use GIS tools as only simple web interfaces are provided to the end user; (b) end users do not have to buy expensive and resource-consuming analysis software (naturally, this does not apply to not service providers); and (c) analysis models can be made accessible to a far broader spectrum of users by exposing them to the web, ranging from private persons who want to know about air quality in their neighbourhood, to urban planners, traffic management officers, and political decision makers.

In conclusion it can be stated that the substantial benefit of the Live Geography approach is that the developed infrastructure is applicable to a wide variety of cross-domain use cases due to its high degree of interoperability, modularity and flexibility, as shown in sub-chapter 8.6.

8.5 Implementation of the Live Geography Approach

The implementation of the Live Geography approach comprises tailor-made sensing devices, a sensor fusion mechanism, a use case specific interpolation model, an automatic server-based analysis component, and a complex event processing and alerting component. All these stand-alone modules are chained together using open standards as described in sub-chapter 8.3.

The system has been conceived in such a modular way that the base platform can be used for a variety of sensor web applications such as environmental monitoring, biometric parameter surveillance, critical infrastructure protection or energy network observation by simply changing the interfaced sensors.

8.5.1 Embedded Geo-Sensing Device

For the measurement device, a special sensing pod for pervasive GIS applications using ubiquitous embedded sensing technologies has been designed. The sensor pod consists of a COTS embedded device, a Gumstix Verdex XM4 platform including an ARM7-based 400MHz processor with 64MB RAM and 16MB flash memory. It runs a customised version of the "Open Embedded" Linux distribution (kernel version 2.6.21) with an overall footprint of <8MB. Additionally to this basic operating system, a GPS module (U-BLOX NEO 4S and LEA-4P) for positioning and several different sensors (e.g. LM92 temperature, NONIN 8000SM oxygen saturation and pulse, or SSM1 radiation sensors) have been attached.

The **embedded software infrastructure** comprises an embedded secure web server supporting HTTP/HTTPS and XMPP, an SQLite database and several daemons, which convert sensor readings before they are provided to the web. The database serves for short-term storage of historic measurements to allow for different error detection procedures and plausibility checks, as well as for unsophisticated trend analysis.

The implementation does not include a dedicated energy saving mode as those mechanisms are already existent and can be integrated in energy critical use cases. Furthermore, sleep mode mechanisms are meanwhile mostly included in new embedded devices currently coming into the market.

8.5.2 Standardised Service Infrastructure

For standardised data retrieval, a service has been created, which implements the OGC Sensor Observation Service (SOS), SensorML and Observations and Measurements (O&M) standards in an application-specific way. Like this,

measurement data are served via HTTP/HTTPS over UMTS in the standardised XML-based O&M format over the SOS service interface. SensorML serves for describing the whole sensor platform as well as in future implementations to control the sensor via the OGC Sensor Planning Service (SPS), e.g. to dynamically adjust measurement cycles.

8.5.3 Real-time Data Integration – Sensor Fusion

For the real-time data integration (sensor fusion) component, a data store extension to the open-source product GeoServer has been developed. This plug-in enables the direct integration of OGC SOS responses (in O&M [Observations and Measurements] format) into GeoServer and their conversion to OGC-conformal service messages on-the-fly.

The main advantage of this approach is that sensor data are made available through a variety of established geo-standardised interfaces (OGC's WFS, WMS, WCS), which offer a variety of output formats such as OGC Geographic Markup Language (GML), OGC Keyhole Markup Language (KML), geoRSS, geoJSON, Scalable Vector Graphics (SVG) or PDF.

Technically speaking, certain input parameters (coordinate reference system, unit conversions, data structures etc.) are interpreted during the transformation procedure from O&M input to WFS output. In practice, this means that the O&M XML structure is converted into well-established standardised data formats as mentioned above.

The innovation in comparison to previous data integration approaches is that the conversion of the data structure from SOS responses to various WFS, WMS and WCS output formats is performed on-the-fly. In contrast, conventional methods typically use the laborious interim step of storing data in a temporary database, which has two distinct disadvantages. First, it adds another component to the overall workflow, which likely causes severe performance losings; secondly, it creates a single central point of failure making the system vulnerable in case of technical malfunctions.

8.5.4 Complex Event Processing and Alerting

Furthermore, a Complex Event Processing (CEP) module coupled with a standards-based alerting module (OGC Sensor Alert Service [SAS]) has been implemented for sending events in a push-based fashion. Within the Live Geography concept, CEP is used for **detecting patterns in measurement data and for creating complex events** accordingly. These can be related to time (temporal validity), space (e.g.

geo-fencing with geographic “intersect”, “overlap” or “join” operations) or measurement parameters (e.g. threshold exceedances). Furthermore, CEP is used for quality assurance in measurement data.

In the actual implementation, SAS is an asynchronous service connecting a sensor in a network to an observation client. In order to receive alerts, a client subscribes to the SAS. If the defined rules apply, a pre-defined alert is sent to the client via the Extensible Messaging and Presence Protocol (XMPP).

This CEP implementation induces a certain degree of intelligence into the embedded device with by providing the possibility to assess measurement data directly on the embedded device on-site. However, these CEP rules feature no self-organising and machine-learning character in this implementation, which is out of scope for this thesis.

8.5.5 Web-based Data Analysis

For the analysis and visualisation modules, several components using ESRI’s ArcGIS software suite have been developed. ArcGIS – including ArcGIS Server – has been chosen as processing engine because reliable, functionally comprehensive and well-established open-source GIS tools are not available yet, especially regarding web-based data processing.

The first analysis component consists of a live sensor extension for ArcMap and ArcScene, which allows for the real-time integration of real-time measurements while having all analysis mechanisms of the ArcToolbox available for data processing. Furthermore, an adapter has been created to integrate sensor data into ESRI’s Tracking Analyst. This enables time-series visualisation and analysis of spatio-temporal patterns. The third analysis component consists of a web-based processing module using ArcGIS Server technology. More technical details of implemented end applications can be found in sub-chapter 8.6.

8.6 Live Geography Portability – Implemented End Applications

At present, real-world sensor network applications are rare and mostly serve a single purpose, which limits broader usage of measurement data. This thesis shows that flexible and portable measurement infrastructures have substantial advantages over closed monolithic systems.

As stated above, monolithic single-purpose sensor networks limit the broad usage of sensor data in cross-domain applications, restrict automated workflows as M2M-communication, and prohibit the achievement of the long-term vision of a “**digital**

skin for the earth" (Gross, 1999), which is comprised of innumerable heterogeneous sensors, discoverable and accessible over the internet.

Thus, researchers in the field of geo-sensor networks have realised the needs of the times and are moving away from pure miniaturisation efforts creating sensor pods, which function exclusively in lab environments. Recent research tends towards standardised **interoperable sensing devices** enabling the establishment of portable domain-independent measurement infrastructures.

This vision of fully integrated and interoperable sensing workflows fosters the awareness for the benefits of open measurement systems. This is especially important for critical tasks in urban environments such as emergency management, environmental monitoring or real-time traffic planning, which are not only relevant to sensor network operators, but also to city management and to citizens.

To demonstrate the Live Geography approach's portability, the following sub-chapters cover four concrete real-world implementations in different application fields. This is to show that the approach is highly portable, interoperable and flexible in terms of trans-domain usage and integration of heterogeneous data sources. This again builds the basis for the deployment of an overarching monitoring infrastructure for solving real-time analysis questions across a variety of research and service areas.

8.6.1 Live Pollutant Monitoring for Public Health

The *Common Scents* project focuses on the use of real-time pollutant monitoring data in the public health sector. As Zardini (2006) states, "we have renounced the utopian idea of a socially, politically, and economically perfect city, but not the promise of a perfectly clean and sanitised environment with pure air for breathing."

Thus, the goal of the project, which is a concerted effort with the MIT SENSEable City Lab, the Research Studio iSPACE, the Harvard University Sensor Networks Lab, the City of Cambridge's Public Health Department, and BBN Technologies, is to provide fine-grained air quality information layers in near real-time. To achieve this vision, the CitySense sensor testbed (Murty et al., 2008) is utilised, measuring CO₂ concentrations together with environmental parameters such as wind speed, wind direction, air temperature, and relative humidity.

The system technologically builds upon the Live Geography approach presented within this thesis. Following principles of Service Oriented Architectures (SOA) and sustainable infrastructure development, components for data collection, processing and information provision for the use case of public health have been conceived.

The empirical project goal was the establishment of an interoperable open standards based infrastructure developed for providing fine-grained air quality data, in order to allow users to assess environmental conditions instantaneously and intuitively. The longer-term aim is to provide citizens with unknown up-to-date information layers in order to support their short-term decisions in real-time. It shall be stated that the term “real-time” in this case is not defined by a pre-set numerical time constant, but more by qualitative expressions such as “immediately” or “ad-hoc”, i.e. information layers have to be created in a timely manner to serve application-specific purposes.

The actual implementation shown in Figure 36 allows for correlating temporal measurement data fluctuation to traffic density, weather conditions or day-time related differences in a very flexible way. The lower left part of Figure 36 shows the temporal gradient of the measurement values (recorded in Cambridge MA, USA), which have been integrated in the standardised O&M format. Running the time series dynamically changes symbologies in the map on the right side accordingly in a dynamic manner.

Preliminary findings show that CO₂ is characterised by very high temporal and spatial fluctuations, which are induced by a variety of factors including temperature variability, time during the day, traffic emergence or “plant respiration” – the fact that plants release major amounts of CO₂ over night. However, the detailed interplay of these parameters still has to be investigated in the next phase of the Common Scents project.




Figure 36: Time Series Visualisation of Pollutant Measurements.

8.6.2 Fine-grained Air Temperature Variations

Another Live Geography implementation has been done in the course of the *Real-time Geo-awareness* project in a cooperative effort with the Research Studio iSPACE and SYNERGIS Informationssysteme GmbH. Apart from the establishment of the technical components (sensor devices, data integration and analysis), the project's aim was to create a sensor network for fine-grained temperature variation assessment.

The pervasive deployment of temperature sensors can lead to detection of urban heat islands with a fine spatial resolution. Furthermore, the temperature measurements can be used for correlation with other environmental parameters such as air pollution, ozone or emissions caused by increased traffic emergence. Thus, an essential part of this particular implementation is the alerting functionality, which is achieved by the use of an OGC Sensor Alert Service (SAS), generating alerts according to pre-defined events, i.e. exceedance of pre-defined thresholds. These events are detected by a Complex Event Processing (CEP) engine that also serves for data quality control by identifying measurement outliers and performing other spatio-temporal plausibility controls.

Figure 37 shows the three-dimensional Inverse Distance Weighting (IDW) interpolation result of air temperature values provided by various OGC Sensor Observation Services (SOS).




Figure 37: Interpolation of Temperature Values for the Detection of Urban Heat Islands.

8.6.3 Ubiquitous Biometric Parameter Surveillance

The *geoHealth Monitor* instance of the Live Geography approach responds to the needs of pervasive medical care. The system uses biometric sensors measuring a person's pulse and oxygen saturation in the blood. The project itself has been carried out in cooperation with the Research Studio iSPACE and Salzburg University of Applied Sciences.

The web interface shown in Figure 38 comprises three sections. Firstly, a configuration panel to select a particular sensing device including different measurement parameters such as the update frequency or the number of measurements stored in the history. The middle section presents the temporal history of OGC Sensor Web Enablement conformal sensor data, which allows for intuitive visual assessment of the measured parameters. Finally, the map on the right side of the interface shows the last few positions of the measurement device to keep track of its spatial trace.




Figure 38: Biometric Parameter History with Geographical Location Illustration.

It shall be stated the *geoHealth Monitor* application cannot only be used for patient surveillance, but may also be employed for equipment tracking, control of the food supply chain including the goods' measured quality condition, or for keeping track of a stolen car.

8.6.4 Real-time Air Quality Assessment

GENESIS (GENeric European Sustainable Information Space for environment), an FP7-funded collaborative research project, has two basic aims: 1.) to establish an open and standards-based infrastructure for managing, analysing and providing environmental information, and 2.) to demonstrate the efficiency of the solution through thematic pilots in different areas within environmental pilot deployments for air quality, water quality and associated health impacts.

The Live Geography approach supports the *GENESIS* project as it builds the technological foundation for the thematic pilots by providing mechanisms for measurement data provision (Sensor Observation Service), sensor fusion (GeoServer data store), alerting (SAS and CEP) and server-based data analysis (ArcGIS Server application).

Using a standardised web client, “live” geo-processing can be performed over the internet combining information in complex geo-processing models (e.g. IHI health index) from various standardised sources like

- OGC SOS live measurements, i.e. sensor fusion data (like temperature, CO₂ and NO₂)
- OGC WMS data for background information (like satellite or orthorectified imagery)
- OGC WFS information layers (like community boundaries, road networks and other geographic feature data)
- OGC WCS data (like MODIS land surface temperature for the IHI index)

Figure 39 illustrates the web interface for live geo-data analysis of environmental information implemented in a kriging process. A special focus in *GENESIS* is on coupling of SAS and CEP including the evaluation of the OGC Sensor Event Service (SES), which is widely seen as the successor of SAS.

In the project, CEP serves for detecting complex patterns in sensor data related to spatial and temporal parameters as well as measurement values. Another emphasis is on integrating legacy GIS systems (ArcGIS Server, GRASS GIS etc.) with the standardised OGC Web Processing Service (WPS) interface to achieve a wholly standardised workflow coupled by a Business Process Execution Language (BPEL) engine.

The Live Geography solution is likely to be integrated into the overall *GENESIS* infrastructure, which finally aims at Europe-wide SDI harmonisation and provision of a complete infrastructure for standardised data access and analysis.




Figure 39: Web-based Live Geo-processing for Environmental Monitoring.

Chapter Summary

The Live Geography approach developed within this thesis stands for the provision of a flexible, modular and interoperable architecture to integrate current sensor data into geospatial analysis systems – particularly targeted to urban environments. Thus, the approach builds the architectural bridge between domain-independent sensor network development and use case specific requirements for end user tailored information output. It aims to complement or replace monolithic single-purpose sensor networks by fully standardised measurement infrastructures.

The challenges for deploying monitoring networks in urban environments are manifold. They range from physical restrictions to privacy issues, knotty political constellations, or questions on data ownership and accessibility.

The developed architecture interlinks components for geo-sensor networks, sensor fusion, event-based data processing (CEP), alerting and web-based data analysis. These modules are chained together by open (geospatial) standards.

To illustrate the Live Geography approach's portability, end applications have been implemented in four different application areas: live pollutant monitoring for public health, assessment of fine-grained air temperature variations for detecting urban heat islands, ubiquitous biometric parameter surveillance for pervasive medical care, and real-time air quality assessment for environmental sustainability.

DISCUSSION AND CONCLUSION

DISCUSSION AND CONCLUSION

9 Discussion and Conclusion

Real-time monitoring of urban dynamics, which is widely unexplored, has recently received a lot of attention due to the fast rise of inexpensive pervasive sensor technologies. Following the comprehensive vision of a “digital skin for planet Earth” by Neil Gross (1999), it can be presumed that geo-sensor web deployments will experience a considerable boost within the coming years, as pervasive sensing has recently become feasible and affordable. This enriches knowledge about our environment with uncharted real-time information layers.

Projections stated that the wide spread distribution of high-speed internet connections would render geographical distance irrelevant. Thus, cities were presumed to drastically decrease in importance as physical and social connections. However in fact, cities play an increasingly important role in socio-technical research. Consequently, cities in their complex structure have re-gained importance in scientific research.

From this process-driven viewpoint, the city is a dynamic multi-layered construct of different processes covering culture, politics, trade, communications infrastructure, finance, technology, energy, traffic, amongst others (s. sub-chapter 3.2). This fact has led to a variety of emerging terms describing the urban environment such as event city, city of bits, cyber city, global city, network city, virtual city or renewable city. This again shows the diverse viewpoints of the urban environment, which yields very different interests.

Thus, when deploying geo-sensor networks for near real-time monitoring purposes in cities, very heterogeneous challenges have to be considered. These range from physical sensor mounting to social and privacy implications. Furthermore, the sensitive urban political landscape has to be accounted for, which might cause unforeseen challenges, as described in sub-chapter 3.3. This thesis discussed these issues and outlined how the *Live Geography* approach can meet resulting requirements.

With these prerequisites and challenges in mind, the *Live Geography* approach has been developed within this thesis. It suggests the establishment of an interoperable, modular and flexible distributed sensing and data analysis infrastructure – as opposed to previous monolithic sensor networks – particularly targeted to urban environments. Thus, it stands for integration of real-time measurement data in a fully standardised infrastructure for urban monitoring applications including web-based data processing. The approach, its potential benefits and its concrete implementation are summarised in the results chapter 8.

9.1 Future Outlook – Complete Live Geography Infrastructure

In its current implementation, the Live Geography approach builds the basis for longer-term developments towards pervasive urban monitoring networks. However, there are a number of future research challenges, as described in sub-chapter 9.1.1.

Moreover, the implemented components presented within this thesis have to be taken further towards applicability in a production environment, which is a matter of source code consolidation and extensive testing. Furthermore, the infrastructure itself has to be extended by additional components to create a fully automated service infrastructure, as illustrated in sub-chapter 9.1.2.

9.1.1 Further Research Directions

Considering that the Live Geography approach is highly interdisciplinary, future research has to be conducted in a variety of areas such as geo-sensor webs, sensor fusion, Complex Event Processing (CEP) for GIS applications, and interoperable web-based geospatial processing.

The main challenge in geo-sensor web research in the coming years will be to harmonise existing networks with upcoming efforts to guarantee optimal data availability for assessing urban dynamics. This implies a shift from developing monolithic single-purpose sensor systems towards creating **interoperable measurement infrastructures**, which requires adequate public awareness and policy frameworks. This in turn allows for the straight-forward use of live sensor data in existing spatial decision support systems.

As the number of geo-sensor network deployments is rapidly growing, a rather pressing challenge will be the integration of huge amounts of measurement data, and their combination with large Earth Observation (EO) and in-situ datasets in **flexible sensor fusion processes**. This requires more research in the area of distributed geo-information architectures and the integration of grid computing into GIS software products.

A tightly connected future research issue is related to CEP. The challenge for CEP is threefold. Firstly, a uniform representation of geospatial CEP rules has to be found to enable machine-to-machine (M2M) communication. Secondly, **coupling CEP engines with GIS software products** will be necessary to establish widely automated geo-workflows, as opposed to current static analysis procedures. Thirdly, the integration of large data sets into CEP engines and connected alerting

technologies needs to be investigated for event detection in a variety of geospatial data, combining terrestrial, airborne and space data.

Next, the integration of **near real-time sensor measurements** in geospatial analysis as additional information layers has to be further investigated. This, considerably increases geospatial analysis quality and accuracy and boosts service-oriented geospatial processing architectures. This is especially important for the distributed infrastructure developed within this thesis. Thus, a special future focus has to be dedicated to the integration of existing tools, models and applications for the different use cases.

Moreover, **standardised web-based geo-processing** is a highly seminal research area, which is still in its infancy. Firstly, several technical challenges have to be tackled from a computer science and computer engineering viewpoint. These comprise efficient resource sharing, load balancing, cloud computing, and high-speed data transfer, particularly for GIS software. Furthermore, standardised processing service interfaces and data structures have to be defined to enable distributed web-based geospatial analysis.

The **OGC Web Processing Service** (WPS) is a promising move in this direction, but it is not clear yet, which functionality it will contain, which data formats it will support, and what process descriptions will exactly look like. The slow and broad development is partly due to different viewpoints within the geospatial community – often driven by individual application contexts. Thus, providing the possibility to integrate user-specific analysis models via the WPS interface will be a crucial point of discussion.

A connected upcoming goal arising from this thesis is to transparently integrate geospatial processing tools into a broader workflow by interfacing it with the standardised WPS service by using **Representational State Transfer (REST) or SOAP technology**. Furthermore, the SOAP and REST based interfaces for geospatial processing have to be extended by a WPS profile compliant interface, which will be developed using the open-source PyWPS technology.

Figure 40 illustrates different approaches for processing standardised sensor fusion data in a service-oriented processing environment. This **federated infrastructure concept** serves the need for ready access to data for near real-time server-based geospatial analysis, web mapping and mobile GIS.




Figure 40: Workflow and Processing Approaches for Sensor Fusion Data.

9.1.2 Fully-fledged Live Geography Infrastructure

As Figure 41 shows, the basic Live Geography service infrastructure as described in chapters 6 and 7 only comprises the essential components for establishing a monitoring system.

The roadmap towards a fully-fledged infrastructure is twofold. First, other SWE standards have to be integrated (Sensor Registry, Sensor Planning Service and Web Notification Service plus an archive database), which is a major development as some standards of the SWE family have not yet been released officially as implementation specifications. Secondly, more complex process models integrating a wider variety of input datasets have to be established and incorporated into the geo-processing infrastructure. This allows for the use of the Live Geography system in sophisticated monitoring and analysis use cases.

This step will be taken in two further short-term developments for urban pollutant dispersion modelling and the calculation of an Integrated Health Index (IHI), as shown in chapter 8.




Figure 41: Fully-fledged Live Geography Infrastructure.

Furthermore, the maturity of the existing components has to be investigated before deployment in a production environment. A brief roadmap for each of the developed components is drawn in the following sub-chapters.

Real-time Data Integration – Sensor Fusion

The sensor fusion component (GeoServer custom data store) for integrating real-time sensor data is ready-to-use. Before integrating it into production environments such as production workflows, several small final refinements (code optimisation, reliability testing, analysis of fitness to application-specific requirements etc.) have to be performed.

Location-aware CEP and Alerting

The SAS and location-aware CEP component exists in a prototypical version. Further discussions will show whether the use of SES or SAS/CEP will be the preferable way to proceed in generating sensor fusion data. The next steps for the developed component will be the integration of several data sources related to a specific use case, the implementation of generic complex geo-rules and the integration of the output alerts into a dedicated workflow using BPEL, e.g. to trigger geo-processing operations according to pre-defined events and alerts.

Web-based Geospatial Data Processing

The implementation for this thesis focused on ArcGIS Server technology, which is particularly important for urban use cases, as public and governmental institutions often employ ArcGIS for geospatial analysis.

Further developments include the transparent integration of ArcGIS Server into a broader workflow by interfacing it with the standardised WPS service by using REST or SOAP technology. Furthermore, the REST and SOAP based interfaces for geospatial processing have to be extended by a WPS profile compliant interface, which will be developed using the open-source PyWPS technology. Moreover, only simple Inverse Distance Weighting (IDW) and kriging interpolation algorithms have been implemented so far. Other processing algorithms have to be implemented according to specific use cases for wider applicability of the system.

9.2 Effects on the Urban Environment

From a socio-political viewpoint, Live Geography is primarily targeted at the information needs of local, regional and national governments as well as international institutions. It enables them to respond to the environmental and social challenges of the city and to learn more about the impacts of urban policies and practices.

Support for Urban Planning and Management

Moreover, it also supports professionals, such as urban and transportation planners, in building or refining their models of urban dynamics. In fact, it can change the work of practitioners that was previously about predicting and accommodating, and which is now becoming more observing and improving.

Indeed, this new ability to render all kinds of “machine readable” environments does not only provide new views on the city and its environment, but also supplies urban and transportation engineers and planners with indicators to evaluate their interventions. For instance, Dan Hill and Duncan Wilson foresee the ability to tune buildings and cities through pre and post occupancy evaluations. They speculate that the future of environmental information will be part of the fabric of buildings (Hill and Wilson, 2008). However, this integration opens all sorts of issues regarding sampling, density, standardisation, quality control, power control, officiality of data, and update frequency.

A complete picture might be hard to achieve with incomplete environmental data patched together by data mining, filtering and visualisation algorithms. Monitoring

in the urban context is limited to classic technical issues related to data resolution and heterogeneity. Even mobile sensors do not yet provide high-density sampling coverage over a wide area, limiting research to sense what is technically possible and to sense with economical and social constraints. One set of solutions rely on the calibration of mathematical models with only a few sensor nodes and complementing data sources to create a set of spatial indicators.

Another approach aims at revealing instead of hiding the incompleteness of the data. Visualising the uncertainty of spatial data is a recurrent theme in cartography and information visualisation (MacEachren et al., 2005). These visualisation techniques present data in such a manner that users are made aware of the degree of uncertainty so as to allow for more informed analyses and decisions. It is a strategy to promote the user appropriation of the information with an awareness of its limitations (Chalmers and Galani, 2004). Without these strategies to handle the fluctuating quality of the data, their partial coverage could impact people's perception of the environment, by providing a quasi-objective but inaccurate angle of the content, and potentially "negatively" influencing their behaviour.

Citizens as Sensors

A way to improve the coverage of environmental data is to alter the current model, whereby civic government would act as sole data gatherer and decision maker by empowering everyday citizens to monitor the environment with sensor-enabled mobile devices (Goodchild, 2007).

Recently providers of geographic and urban data have learned the value of people-centric sensing to improve their services and from the activities of their customers. For instance the body of knowledge on a city's road conditions and real-time road traffic network information thrive on the crowd-sourcing of geo-data by collecting data, which the owners of navigation systems and mobile phones generate. Similarly, the users of Google MyMaps have contributed, without their awareness, to the production the massive database necessary for the development of the location-based version of the application.

However, this people-centric approach to gather data raise legitimate privacy concerns. These issues can be handled with a mix of policy definition, local processing, verification and privacy preserving data mining techniques (Abdelzaher et al., 2007). These technical solutions necessitate a richer discussion beyond the academic domain on these observing technologies' social implications.

Similar crowd-sourcing strategies have been considered for urban monitoring with individuals acting as sensor nodes and coming together with other people in order

to form sensor networks. Several research projects explore a wide range of novel physical sensors attached to mobile devices empowering everyday non-experts to collect and share air quality data measured with sensor-enabled mobile devices. For instance, Ergo (Paulos, 2008) is a simple SMS system that allows anyone with a mobile phone to quickly and easily explore, query, and learn about local air quality on-the-go with their mobile phone.

With these tools, citizens augment their role, becoming agents of change by uncovering, visualising, and sharing real-time air quality measurements from their own everyday urban life. This "citizen science" approach (Paulos et al., 2008) creates value-added information for researchers induced by data generated by people going on their daily life, often based on explicit and participatory sensing actions. By turning mobile phones (Nokia Research, 2008), watches (La Montre Verte, 2009) or bikes (Campbell et al., 2006) into sensing devices, the researchers hope that public understandings of science and environmental issues will be improved and can have access to larger and more detailed data sets. This access to environmental data of the city also becomes a tool to raise citizens' awareness of the state of the environment.

These data gathering possibilities also imply that we are at the end of the ephemeral; in some ways we will be able to replay the city. In contrast, we are also ahead of conflicts to reveal or hide unwanted evidences, when new data can be used to the detriment of some stakeholders and policy makers. Indeed, the capacity to collect and disseminate reconfigure sensor data influence political networks, focussing on monitoring information as products or objects that can be used for future political action.

Therefore, the openness, quality, confidence, and trust in the data will also be subject of debate (e.g. bias to have people record their observations; who gets to report data and who not). This people-centric view of measuring, sharing, and discussing urban environments might increase agencies' and decision makers' understanding of a community's claims, thereby potentially increasing public trust in the information provided by a real-time geo-awareness approach.

Selective Induction of Behavioural Change

This raises a real challenge: how can we encourage, promote and motivate environmentally sustainable behaviours on the basis of Live Geography information? Recent work (Calabrese et al., 2008) shows how real-time information about cities and their patterns of use, visualised in new ways, and made available locally on-demand in ways that people can act upon, may make an important

contribution to sustainability. However, the communication of data collected from pervasive sensor networks may not trigger sufficient motivation for people to change their habits towards a more environmentally sustainable lifestyle.

Indeed, this objective of improving a city's environmental sustainability calls for behaviour modification. It can be induced by intervening in moments of local decision making and by providing people with new rewards and new motivations for desirable behaviours (Fogg, 2003). These kinds of strategies have been common, for instance, in health and fitness applications.

However, when we think about persuasion in the real of environment sustainability, we might want to persuade people of the ways, in which their interests are aligned with those of others (Dourish, 2008). Therefore, this process of alignment and mobilisation, by which one can start to find one's own interests as being congruent with those of others will be critical in the success of these strategies based on Live Geography.

Sensitising People to Their Context

All of the above examples consider the use of sensed data to satisfy public or private requirements of some sort. However sensor data are just indicators, e.g. for the effects of pollutants on people or structures, and there is much disagreement on the public information regulations and whether they are effective. One potential alternative is the direct sensing of structural, or even human, impacts. In other words, people need to be sensitised to their location and their environment, and this information will have to be presented within different contexts, such as safety, public health, energy consumption or sustainability in order to open new possibilities for exploring the city.

In sum, feedback of "sensed" data to the public – either directly or after processing – can potentially change people's perception of the city itself. For example, weather or pollution prediction on a very local scale in space and time is reasonably feasible. Continuous availability of this information could lead to a change in people's individual behaviour by giving real-time short-term decision support. The Live Geography approach can potentially play a significant role in achieving this seminal vision.

9.3 Changing the Perception of the City

Generally speaking, fine-grained urban sensing greatly enhances our knowledge of the city by adding objective and non-visible data layers in real-time. This seems to be a very desirable state because more accurate data about local air temperature,

atmospheric humidity, gaseous and particulate air pollution, and traffic emissions can positively influence areas such as public health, traffic management or emergency response.

Measurement Data as Surrogates for Direct Impact on Humans

Apart from this information enrichment, accurate sensor measurements also have a much broader influence: considering for example that “air quality” is only a surrogate for the effects of pollutants on humans makes a fine-grained air quality map a very sensitive information layer, as discussed in sub-chapter 3.3.

In terms of public health, the primary impact of the Live Geography approach can be a change in the perception of the measured parameters among the population. Seeing greater differences in air quality between one neighbourhood and another will present a more complex picture of the urban environment that is not based on broad well-known generalities like “basin-wide” air quality data, but it will show that air quality can be very locally heterogeneous.

Fine-grained monitoring information such as air quality has a direct influence on people’s health and life expectation; such information could potentially induce far-reaching changes in different areas relating to public health such as health care provision, the personal insurance sector, the real-estate market or policy making in the area of individual traffic.

One effect could be that greater access to data shows certain residents at greater risk than others, which will result in essential questions like: Does the disparity in air quality across the city really reflect increased lifetime risk of morbidity or premature mortality? What can the city do to mitigate the impact of these exposures in area that are more impacted? Will decisions be made using this sort of data on the basis of a disciplined risk assessment process or will perception and emotions drive city decisions on planning and land-use? A more practical question could be: Would a much-needed school not be built because of proximity to a large roadway showing some elevation of hazardous air pollutants, but an increase that would not be considered to be statistically significant in the overall health of children who would attend?

Discussion on Environmental Justice

From a quantitative viewpoint, the air quality data measured with the Common Scents implementation of the Live Geography infrastructure will probably show that some areas of the city are “cleaner” by a factor of two, which is indicated by recent preliminary NO_x studies conducted by the City of Cambridge’s Public Health

Department. Thus, the measurement data will certainly contribute to the discussion now active on “environmental justice” by providing highly specific data that indicates, in a vivid fashion, the disparities in environmental burden that have mostly been acknowledged in the abstract up to this point.

Also, inducing public health related sensor data can have a significant impact on urban planning in general. Possibly, decisions about locations of housing developments, daycares, schools or recreational areas can be supported. Furthermore, data can be added to environmental justice assessments, which may result in a stronger case (greater disparity by social class) or a weaker case (less disparity by social class).

Also, it may help our cities adopt better and more progressive transportation policies, giving greater emphasis to walking paths, bicycle routes and alternative-fuel vehicles by raising increased awareness for our physical environment. In other words, it will help being citizens to appreciate the interrelated nature of cities as ecosystems, generally raising consciousness among residents and decision makers about the organic nature of successful urban ecologies.

Individual Real-time Decision Support

Regarding the individual it has to be considered that – in the future of pervasive sensing – pedestrians, cyclists or car drivers may likely be carrying an ultra-small sensing array within their personal GIS system that will provide them with information on the environment they are experiencing and make this automatically available (after some manipulation) to others via the pervasive sensing network. In effect, such a system allows people to take short-term informed decisions, as opposed to current approaches, which can give people decision support at best at coarse scale and for the long-term.

Naturally, this raises issues of privacy, data ownership, accessibility and integrity, which have to be thoroughly tackled. That this should preferably happen by overarching legal frameworks instead of as the case arises, separately for every project.

Finally, concerns might arise about how such fine-grained data and information would be used to eliminate even more mystery from our lives. So far, the blurriness of our impressions about the world we live in, especially the ambient mystique associated with our cities, is part of our lives and our perception of the world. Introducing new information layers could result in the feeling that this mystique might be exposed or objectified to the point that there is nothing left to envision, especially if it is used for commercial gain.

9.4 Conclusion

Ubiquitous and continuous monitoring is a multi-dimensional challenge, and this is particularly true in the urban context. This thesis shows, which issues have to be considered for pervasive monitoring systems in the city, and outlines how the *Live Geography* approach can meet these requirements.

Live Geography Approach

The approach presented in this thesis stands for the combination of live sensor data with historic data sources in an open standards based framework using server-side processing mechanisms. In other words, it tries to establish an interoperable, modular and flexible sensing and data analysis infrastructure, as opposed to previous monolithic sensor networks.

Its basic aim is twofold: first, to complement or replace monolithic measurement systems with an open standards-based infrastructure; second, to automate geospatial analysis processes, using new server-based processing tools by providing a wholly standardised workflow. Since interoperable open systems are not trivial to implement, this motivation has to be initiated through legal directives like INSPIRE, GMES, SEIS and GEOSS.

Portability – Independence of Specific Application Domains

To prove the system's portability, the same underlying framework has been deployed in different application areas: public health (pollutant concentrations and associated health effects), air quality monitoring (real-time computation of a multi-dimensional air quality indicator), patient surveillance (monitoring of biometric parameters), and urban heat island detection (air temperature variation assessment), as described in chapter 8.

The next practical realisation will be an urban "quality of life" monitoring system in order to gain a city-wide picture of spatial and temporal variations of environmental parameters (particulate matter, pollen concentrations or CO pollution). Their integration with static geospatial data (census, traffic emergence, living vs. working places etc.) will maximise significance for local and regional governments, as well as for citizens by applying complex geospatial algorithms such as co-kriging to reveal unknown causal correlations between spatial parameters. Other scheduled implementations comprise radiation monitoring and water quality assessment.

As for enhanced usability on the end user side, a mobile GIS application is currently under development, which allows the user to interact with the system while offering

various geographic analysis tools. Hence, people can not only visualise data, but they can generate personalised information layers by using pre-defined geospatial algorithms in a one-click manner. Apart from that, the mobile device also allows for feeding back data into the system, e.g. by digitising geographic real-world features.

Potential Influence on the Urban Environment

As the described implementations are operated in close cooperation with local or regional governments and thematic actors, it is likely that the developed approach has the potential to raise awareness of ubiquitous sensing systems and perhaps trigger profound rethinking process in collaboration and cooperation efforts between different authorities in the city. Thus, the approach can probably respond to dedicated needs of the city management. Therefore, the longer-term goal is to enhance people's awareness of their environment by adding previously unseen information layers and thus offering real-time decision support to influence their short-term actions.

From a qualitative viewpoint, this thesis also elucidated that the impact of pervasive sensing in the city has to be carefully assessed. It was found that providing very fine-grained information layers might on the one hand be a powerful decision support instrument, but on the other hand too detailed environmental information might also have negative effects. As "air quality" is just a surrogate for more personal impacts such as life expectation or respiration diseases, this information could yield a very broad impact in various kinds of areas such as housing markets, the insurance sector or urban planning in general.

Concluding, it shall be stated that the trend towards extensive availability of measurement data requires a paradigm shift in the global GIS market and its applications for urban environmental monitoring. This applies especially to open data accessibility and intensified collaboration efforts. To achieve far-reaching adoption, the establishment of ubiquitous sensing infrastructures will require a long process of sensitising individuals to their spatial and social contexts, and to the connection of local environmental questions to well-known urban issues such as public safety, energy efficiency, or social interaction.

In effect, creating a meaningful context around densely available sensor data and distributing specific information layers makes the environment more understandable to the city management, to the citizens and to researchers by "making the abstract real", i.e. by revealing hidden connections in real-time.

Chapter Summary

When deploying geo-sensor networks for near real-time monitoring purposes in cities, very heterogeneous challenges have to be considered. These range from physical sensor mounting to social and privacy implications to the sensitive urban political landscape has to be accounted for, which might cause unforeseen challenges.

With these prerequisites and challenges of urban monitoring in mind, the Live Geography approach has been developed within this thesis. It suggests the establishment of an interoperable, modular and flexible distributed sensing and data analysis architecture – as opposed to previous monolithic sensor networks – particularly targeted to urban environments. Thus, it stands for integration of real-time measurement data in a fully standardised infrastructure for urban monitoring applications including web-based data processing.

The architecture includes components for geo-sensor networks, sensor fusion, event-based data processing (CEP), alerting and web-based data analysis. These modules are chained together by open (geospatial) standards.

To demonstrate the Live Geography approach's transferability, four concrete real-world implementations in four different application fields have been presented in this chapter. This is to show that the approach is highly portable in terms of trans-domain usage and integration of heterogeneous data sources. This again builds the basis for the deployment of an overarching monitoring infrastructure for solving real-time analysis questions across a variety of research and service areas.

From a qualitative viewpoint, this thesis also elucidated that the impact of pervasive sensing in the city has to be carefully assessed. It was found that e.g. providing very fine-grained information layers might on the one hand be a powerful decision support instrument, but on the other hand too detailed environmental information might also have negative effects. As "air quality" is just a surrogate for more personal impacts such as life expectation or respiration diseases, this information could yield a very broad impact in various kinds of areas such as housing markets, the insurance sector or urban planning in general.

EPILOGUE

ЕПИЛОГ

References

- 52North (2009) 52n WPS Geoprocessing Community. <http://52north.org>, 2009. (24 September 2009)
- Abdelzaher, T., Anokwa, Y., Boda, P., Burke, J., Estrin, D., Guibas, L., Kansal, A., Madden, S. and Reich, J. (2007) Mobiscopes for Human Spaces. *IEEE Pervasive Computing*, 6(2), pp. 20–29, 2007.
- Accenture Technology Labs (2004) Remote Sensor Network: Accenture Prototype Helps Pickberry Vineyard Improve Crop Management. <http://www.accenture.com>, 2004. (28 April 2005)
- Alberts, D.S., Garstka, J.J. and Stein, F.P. (Eds.) (1999) *Network Centric Warfare: Developing and Leveraging Information Superiority*. 2nd Edition (Revised), Command and Control Research Program (CCRP) Press, Washington, D.C., USA, September 1999.
- Anderson, D.M. (2004) *Build-to-Order & Mass Customization – The Ultimate Supply Chain Management and Lean Production Strategy for Low-Cost On-Demand Production without Forecasts or Inventory*. 520 pp., CIM Press, Cambria, CA, USA, 2004.
- Arasu, A., Babu, S. and Widom, J. (2003) *The CQL Continuous Query Language: Semantic Foundations and Query Execution*. Technical Report. Stanford, <http://ilpubs.stanford.edu:8090/758/1/2003-67.pdf>, October 2003. (29 August 2009)
- Baker, F. and Carpenter, B. (2005) Structure of an International Emergency Alert System. <http://www.ietf.org>, IETF Internet-Draft, 10 January 2005. (2 May 2005)
- Balazinska, M., Deshpande, A., Franklin, M.J., Gibbons, P.B., Gray, J., Hansen, M., Liebhold, M., Nath, S., Szalay, A. and Tao, V. (2007) Data Management in the Worldwide Sensor Web. *IEEE Pervasive Computing*, 6(2), pp. 30–40, Apr-Jun 2007.

- Bass, T. (2007) What is Complex Event Processing?.
<http://www.thecepblog.com/what-is-complex-event-processing>, 2007. (11 July 2009)
- Baumann, P. (2009) Web Coverage Processing Service (WCPS) Language Interface Standard. <http://www.opengeospatial.org>, OpenGIS Interface Standard OGC 08-068r2, Version 1.0.0, 25 March 2009. (28 August 2009)
- Biron, P. und Malhotra, A. (Eds.) (2004) XML Schema Part 2: Datatypes Second Edition. <http://www.w3.org/TR/xmlschema-2>, W3C Recommendation Paper, Second Edition, 28 October 2004. (23 May 2005)
- Borsch, S. (2009) Distance is Dying.
http://blogs.scholastic.com/accelerating_change/2009/07/distance-is-dying.html, 7 July 2009. (21 August 2009)
- Botts, M. (Ed.) (2004) Sensor Model Language (SensorML) for In-situ and Remote Sensors. <http://www.opengeospatial.org>, OpenGIS Consortium Implementation Specification 04-019r2, Version 1.0.0 beta, 1 November 2004. (8 April 2005)
- Botts, M., Robin, A., Davidson, J. and Simonis, I. (Eds.) (2006) OpenGIS Sensor Web Enablement Architecture Document. <http://www.opengeospatial.org>, OpenGIS Interoperability Project Report 06-021r1, Version 1.0, 4 March 2006. (6 December 2008)
- Botts, M. (Ed.) (2007) OpenGIS Sensor Model Language (SensorML) Implementation Specification. <http://www.opengeospatial.org>, OpenGIS Implementation Specification OGC 07-000, Version 1.0.0 beta, 17 July 2007. (8 August 2009)
- Botts, M., Percivall, G., Reed, C. and Davidson, J. (Eds.) (2007) OGC Sensor Web Enablement: Overview And High Level Architecture.
<http://www.opengeospatial.org>, OpenGIS White Paper OGC 07-165, Version 3, 28 December 2007. (17 August 2009)
- Cairncross, F. (1997) The Death of Distance: How the Communications Revolution Will Change Our Lives. Harvard Business School Press, Boston, MA, USA, 1997.

- Calabrese, F., Kloeckl, K. and Ratti, C. (2008) WikiCity: Real-time Location-sensitive Tools for the City. In: Handbook of Research on Urban Informatics: The Practice and Promise of the Real-Time City. 506 pp., ISBN 978-1605661520, IGI Global, Hershey, PA, USA, 2008.
- Campbell, A.T., Eisenman, S.B., Lane, N.D., Miluzzo, E. and Peterson, R.A. (2006) People-centric Urban Sensing. In: Proceedings of the 2nd Annual International Workshop on Wireless Internet – WICON'06, ACM, New York, NY, USA, 2006.
- Castells, M. (1996) The Rise of Network Society. In: Castells, Manuel (1996) The Information Age – Economy, Society and Culture. vol. 1, Blackwell, Oxford, UK, 1996.
- Center for Coastal and Land-Margin Research (2009) CORIE.
<http://www.ccalmr.org/CORIE>, June 2009. (14 July 2009)
- Cepicky, J. (2009) Welcome to PyWPS. <http://pywps.wald.intevation.org>, June 2009. (8 October 2009)
- Chalmers, M. and Galani, A. (2004) Seamless Interweaving: Heterogeneity in the Theory and Design of Interactive Systems. In: Proceedings of the 2004 Conference on Designing Interactive Systems – DIS'04, pp. 243–252, ACM Press, New York, NY, USA.
- Chan, H. and Perrig, A. (2003) Security and Privacy in Sensor Networks. IEEE Computer Magazine, 36(10), pp. 103-105, October 2003.
- Choi, J., Yoo, S., Park, H., and Chun, J. (2005) Seamless Real-time Clinical Data Integration for Mobile Clinical Information System. In: Proceedings of the 21st International Conference on Data Engineering Workshops (ICDEW), Tokyo, Japan, 5-8 April, 2005.
- Chong, C.-Y. and Kumar, S.P. (2003) Sensor Networks: Evolution, Opportunities, and Challenges. Proceedings of the IEEE, 91(8), pp. 1247-1256, August 2003.
- Churchill, P.N. and Sharman, M.J. (1995) The Centre for Earth Observation (CEO): An Initiative of the European Commission to Promote the Use and Availability of Earth Observation Data. International Geoscience and Remote Sensing

- Symposium, IGARSS '95. Quantitative Remote Sensing for Science and Applications, vol. 2, pp. 1551-1553, Florence, Italy, July 1995.
- Coimbra Delicato, F., Pires, P.F., Pirmez, L. and Rust da Costa Carmo, L.F. (2003) A Flexible Web Service based Architecture for Wireless Sensor Networks. IEEE MWN 2003 – Workshop on Mobile and Wireless Networks, May 2003.
- CollabNet Inc. (2009a) Spring Plug-in for OGC's Web Processing Service (WPS) Interface. <http://wpsint.tigris.org>, 2009. (14 July 2009)
- CollabNet Inc. (2009b) Open-esb: IEP – Open Source Complex Event Processing (CEP) and Event Stream Processing (ESP) Engine. <https://open-esb.dev.java.net/IEPSE.html>, August 2009. (27 August 2009)
- Collins, J. (2004) NASA Creates Thinking RF Sensors. RFID Journal, October 2004.
- Commission of the European Communities (2008) Towards a Shared Environmental Information System (SEIS). Communication from the Commission to the Council, the European Parliament, the European Economic and Social Committee and the Committee of the Regions, <http://ec.europa.eu/environment/seis>, Brussels, 1 February 2008. (10 September 2009)
- Commission of the European Communities (2009a) INSPIRE Directive. <http://inspire.jrc.ec.europa.eu>, August 2009. (29 August 2009)
- Commission of the European Communities (2009b) European Commission – GMES Home Page. <http://ec.europa.eu/gmes>, June 2009. (17 June 2009)
- Computational Geoscience Group (2009) SEE Grid. <http://www.seegrid.csiro.au>, 18 February 2009. (6 September 2009)
- Cox, S. (Ed.) (2007) Observations and Measurements. <http://www.opengeospatial.org>, OpenGIS Implementation Standard OGC 07-022r1, Version 1.0, 8 December 2007. (12 May 2009)

- de la Beaujardiere, J. (Ed.) (2006) OpenGIS Web Map Server Implementation Specification. <http://www.opengeospatial.org>, OpenGIS Implementation Specification OGC 06-042, Version 1.3.0, 15 March 2006. (12 June 2009)
- Delin, K.A., Harvey, R.P., Chabot, N.A., Jackson, S.P., Adams, M., Johnson, D.W. and Britton, J.T. (2003) Sensor Web in Antarctica: Developing an Intelligent, Autonomous Platform for Locating Biological Flourishes in Cryogenic Environments. 34th Lunar and Planetary Science Conference, Houston, TX, USA, March 2003.
- Delin, K.A., Jackson, S.P., Johnson, D.W., Burleigh, S.C., Woodrow, R.R., McAuley, J.M., Dohm, J.M., Ip, F., Ferré, T.P., Rucker, D.F. and Baker, V.R. (2005) Environmental Studies with the Sensor Web: Principles and Practice. *Sensors*, vol. 5, pp. 103-117, 2005.
- Delin, K.A. (Web Curator) and Woodrow, R. (Maintenance Responsible) (2005) NASA/JPL Sensor Webs Project. <http://sensorwebs.jpl.nasa.gov>, 13 April 2005. (26 April 2005)
- Dierig, S., Lachmund, J., and Mendelsohn, A. (2000) Science and the City. <http://vlp.mpiwg-berlin.mpg.de>, Workshop, Max Planck Institute for the History of Science, Berlin, Germany, 1-3 December 2000. (10 September 2009)
- DIN German Institute for Standardization (2000) Economic Benefits of Standardization – Summary of Results, Final Report and Practical Examples. 40 pp., ISBN 3-410-14860-4, Beuth Verlag, Berlin, Germany, August 2000.
- Douglas, J., Usländer, T., Schimak, G., Esteban, J.F. and Denzer, R. (2008) An Open Distributed Architecture for Sensor Networks for Risk Management. *Sensors*, vol. 8, pp. 1755-1773, 2008.
- Dourish, P. (2008) Points of Persuasion: Strategic Essentialism and Environmental Sustainability. In: Proceedings of the 6th International Conference on Pervasive Computing, Sydney, Australia, 19-22 May 2008.
- Droege, P. (2006) Renewable City – A Comprehensive Guide to an Urban Revolution. 309 pp., Wiley-Academy a division of John Wiley & Sons Ltd, ISBN 978-0-470-01925-2, Chichester, England, 2006.

- Echterhoff, J. and Everding, T. (2008) OpenGIS Sensor Event Service Implementation Specification. <http://www.opengeospatial.org>, OpenGIS Discussion Paper OGC 08-133, Version 0.0.1, 27 August 2008. (20 August 2009)
- Ecole des Mines de Paris (2009) The Data Fusion Server. <http://www.data-fusion.org>, August 2009. (21 August 2009)
- EsperTech Inc. (2009) Esper – Complex Event Processing. <http://esper.codehaus.org>, August 2009. (24 August 2009)
- ESRI (2009) What is Geo-processing?. http://edndoc.esri.com/arcobjects/9.2/NET/shared/geoprocessing/geoprocessing/what_is_geoprocessing_qst_.htm, ESRI Developer Network Documentation, August 2009. (27 August 2009)
- Fischer, F. (2007) GIScience – Connecting Real and Virtual Worlds. Presentation at the EGEA Annual Congress 2007, 17-22 September 2007, Ustroń, Poland.
- Fogg, B.J. (2003) Persuasive Technology: Using Computers to Change What We Think and Do. 312 pp., ISBN 978-1558606432, Morgan Kaufmann, Amsterdam, The Neatherlands, 2003.
- Friedmann, J. (1986) The World City Hypothesis. *Development and Change*, 17(1), pp. 69-83, 1995.
- General Services Administration (1996) Federal Standard 1037C. Telecommunications: Glossary of Telecommunication Terms, <http://www.its.bldrdoc.gov>, 7 August 1996. (11 September 2009)
- Gibbons, P.B., Karp, B., Ke, Y., Nath, S. and Seshan, S. (2003) IrisNet: An Architecture for a Worldwide Sensor Web. *IEEE Pervasive Computing*, Special Issue on Sensor and Actuator Networks, 2(4), pp. 22-33, IEEE Press, October-December 2003.
- Gilder, G. and Peters, T. (1995) City vs. Country: The Impact of Technology on Location. *Forbes ASAP*, 155(5), pp. 56-61, 27 February 1995.

- Goodchild, M.F. (2007) Citizens as Voluntary Sensors: Spatial Data Infrastructure in the World of Web 2.0. *International Journal of Spatial Data Infrastructures Research*, vol. 2, pp. 24-32, 2007.
- Gross, N. (1999) 14: The Earth Will Don an Electronic Skin. <http://www.businessweek.com>, BusinessWeek Online, 30 August 1999. (20 September 2005)
- Hall, P. (1966) *The World Cities*. 3rd Edition, 256 pp., Weidenfeld and Nicolson, London, UK, 1966.
- Hall, P. and Pfeiffer, U. (2000) *URBAN 21 – Expertenbericht zur Zukunft der Städte*. 453 pp., Deutsche Verlagsanstalt, Stuttgart, Germany, 2000.
- Harrie, L. (2004) Using Simultaneous Graphic Generalisation in a System for Real-Time Maps. *Papers of the ICA Workshop on Generalisation and Multiple Representation*, Leicester, August 20-21, 2004.
- Herring, J.R. (Ed.) (2006) *OpenGIS Implementation Specification for Geographic Information – Simple Feature Access – Part 2: SQL Option*. <http://www.opengeospatial.org>, OpenGIS Implementation Specification OGC 06-104r3, Version 1.2.0, 5 October 2006. (20 August 2009)
- Hildebrand, P., Schoeberl, M., Wiscomb, W., Albjerg, M., Mlynczak, M., Raymond, C., Ferraro, R., Miller, T., Miller, R. and Petersen, D. (2003) *The Future Global Earth Observing System: System Requirements and Architecture*. IEEE/Geoscience and Remote Sensing Society IGARSS-03, Toulouse, France, 21-25 July 2003.
- Hill, D. and Wilson, D. (2008) *The New Well-tempered Environment: Tuning Buildings and Cities. Pervasive Persuasive Technology and Environmental Sustainability*, 2008.
- Hipp, R. (2009) SQLite Home Page. <http://www.sqlite.org>, 1 August 2009. (5 September 2009)

- Intel Corporation (2005) Research – Vertical Industry Applications – Manufacturing/Energy - BP Preventative Maintenance Tanker Project. http://www.intel.com/research/vert_manuf_prevmaint.htm, 2005. (2 May 2005)
- IAA Alliance (2005) IAA Alliance. <http://www.iaa-alliance.net>, April 2005. (28 April 2005)
- IBM Corporation (2005) IBM United States. <http://www.ibm.com>, 2009. (25 August 2007)
- International Organization for Standardization (2005a) Geographic Information – Geography Markup Language (GML). <http://www.isotc211.org>, ISO/TC211 Standard, Reference Number 19136, 15 April 2005. (24 May 2005)
- International Organization for Standardization (2005b) Geographic Information – Services. <http://www.iso.org>, 67 pp., 1 February 2005. (24 August 2009)
- Kahn, J.M., Katz, R.H. and Pister, K.S.J. (1999) Next Century Challenges: Mobile Networking for “Smart Dust”. International Conference on Mobile Computing and Networks (MobiCOM), pp. 271–278, Seattle, Washington, August 1999.
- Kaler, C. (Ed.) (2002) Web Services Security (WS-Security). <http://www-106.ibm.com>, Version 1.0, 5 April 2002. (12 May 2005)
- Kansal, A., Nath, S., Liu, J. and Zhao, F. (2007) SenseWeb: An Infrastructure for Shared Sensing. IEEE Multimedia, 14(4), pp. 8-13, October-December 2007.
- King’s College London (2009) The London Air Quality Network. <http://www.londonair.org.uk>, August 2009. (6 September 2009)
- Klappenbach, D., Meissner, A., and Ursa, Y. (2004) Linking Trucks and Central Operations Base – Multimedia Enhancements for the Coldroad System. Second International Conference on Advances in Mobile Multimedia (MOMM2004), Bali, Indonesia, ISBN 3-85403-184-X, pp. 87-91, 22-24 September 2004.
- La Montre Verte (2009) La Montre Verte City Pulse. <http://www.lamontreverte.org>, July 2009. (10 August 2009)

- lat/lon GmbH (2009) deegree Web Processing Service v2.2. <http://deegree.org>, August 2009. (3 October 2009)
- Lehto, L. and Sarjakoski, L.T. (2005) Real-time Generalisation of XML-encoded Spatial Data for the Web and Mobile Devices. *International Journal of Geographical Information Science*, 19(8-9), pp. 957-973, 2005.
- Liang, S., Toa, V. and Croitoru, A. (2004) Sensor Web and Geoswift – an Open Geospatial Sensing Service. In: *Proceedings of the 20th International Society for Photogrammetry and Remote Sensing Congress – Geo-Imagery Bridging Continents*, Istanbul, Turkey, 12-23 July 2004.
- Lindsey, S. and Raghavendra, C. (2002) PEGASIS: Power-Efficient Gathering in Sensor Information Systems. In: *Proceedings of the IEEE Aerospace Conference*, 3(9-16), pp. 1125-1130, March 2002.
- Llinas, J., Bowman, C., Rogova, G., Steinberg, A., Waltz, E. and White, F. (2004) Revisiting the JDL Data Fusion Model II. In: *Proceedings of the Seventh International Conference on Information Fusion*, Stockholm, Sweden, 28 June – 1 July 2004.
- Mabrouk, M. (Ed.) (2008) OpenGIS Location Services (OpenLS): Core Services. <http://www.opengeospatial.org>, OpenGIS Interface Standard OGC 07-074, Version 1.2, 9 September 2008. (9 September 2009)
- MacEachren, A.M., Robinson, A., Hopper, S., Gardner, S., Murray, R., Gahegan, M. and Hetzler, E. (2005) Visualizing Geospatial Information Uncertainty: What We Know and What We Need to Know. *Cartography and Geographic Information Science*, 32(3), pp. 139-160, July 2005.
- Mhatre, V. and Rosenberg, C. (2004) Homogeneous vs Heterogeneous Clustered Sensor Networks: A Comparative Study. In: *Proceedings of IEEE International Conference on Communications*, Paris, France, June 2004.
- Minami, M., Saruwatari, S., Kashima, T., Morito, T., Morikawa, H. and Aoyama, T. (2004) Implementation-based Approach for Designing Practical Sensor Network

- Systems. In: Proceedings of the 11th Asia-Pacific Software Engineering Conference (APSEC '04), pp. 703-710, 30 November – 3 December 2004.
- Mitchell, William J. (1996) *City of Bits – Space, Place, and the Infobahn*. 225 pp., MIT Press, Cambridge MA, USA, 1996.
- Murty, R., Mainland, G., Rose, I., Chowdhury, A., Gosain, A., Bers, J. and Welsh, M. (2008) *CitySense: A Vision for an Urban-Scale Wireless Networking Testbed*. In: Proceedings of the 2008 IEEE International Conference on Technologies for Homeland Security, Waltham, MA, May 2008.
- Na, A. and Priest, M. (Eds.) (2007) *Sensor Observation Service*. <http://www.opengeospatial.org>, OpenGIS Implementation Standard OGC 06-009r6, Version 1.0, 26 October 2007. (12 August 2009)
- Nash, E. (2008) *WPS Application Profiles for Generic and Specialised Processes*. In: Proceedings of the GI Days, Münster, Germany, 16-17 June 2008.
- National Oceanic and Atmospheric Administration (2008) *nowCOAST: GIS Mapping Portal to Real-Time Environmental Observations and NOAA Forecasts*. <http://nowcoast.noaa.gov>, September 2008. (15 December 2008)
- Nebert, D. (2002) *OpenGIS Catalog Services Specification*. <http://www.opengeospatial.org>, OpenGIS Consortium Implementation Specification, Version 1.1.1, 13 December 1999. (19 May 2005)
- Netherlands Organization for Scientific Research (2007) *Urban Sciences*. <http://www.urbansciences.eu>, Interdisciplinary Research Programme on Urbanization & Urban culture in The Netherlands, 2007. (26 August 2009)
- Nickerson, B. and Lu, J. (2004) *A Language for Wireless Sensor Webs*. Communication Networks and Services Research Conference (CNSR), pp. 293-300, Fredericton, New Brunswick, Canada, May 2004.
- Niedzwiadek H. (2002) *OpenLS Platform*. <http://www.jlocationsservices.com>, 20 December 2002. (21 April 2005)

- Nokia Research (2008) SensorPlanet. <http://www.sensorplanet.org>, 12 March 2008. (18 August 2009)
- Obraczka, K., Manduchi, R. and Garcia-Luna-Avecas, J.J. (2002) Managing the Information Flow in Visual Sensor Networks. International Symposium on Wireless Personal Multimedia Communications (WMPC), vol. 3, ISBN 0-7803-7442-8, pp. 1177-1181, October 2002.
- Onsrud, H. (Web Administrator) (2009) Global Spatial Data Infrastructure Association. <http://www.gsdi.org>, July 2009. (14 July 2009)
- Open Source Geospatial Foundation (2009) GeoTools The Open Source Java GIS Toolkit. <http://www.geotools.org>, 2009. (24 September 2009)
- OpenGeo (2009) Welcome – GeoServer. <http://geoserver.org>, 14 October 2009. (14 October 2009)
- Oracle Corporation (2009) Oracle Complex Event Processing. <http://www.oracle.com/technology/products/event-driven-architecture/complex-event-processing.html>, August 2009. (28 August 2009)
- Pan, A. and Viña, A. (2004) An Alternative Architecture for Financial Data Integration. Communications of the ACM (CACM), 47(5), pp. 37-40, May 2004.
- Paulos, E. (2008) Urban Atmospheres - Proactive Archeology of Our Urban Landscapes and Emerging Technology. <http://www.urban-atmospheres.net>, 2008. (18 June 2009)
- Paulos, E., Honicky, R. and Hooker, B. (2008) Citizen Science: Enabling Participatory Urbanism. In: Handbook of Research on Urban Informatics: The Practice and Promise of the Real-Time City. 506 pp., ISBN 978-1605661520, IGI Global, Hershey, PA, USA, 2008.
- Paulsen, H. (2008) PermaSensorGIS – Real-time Permafrost Data. Geoconnexion International Magazine, 02/2008, pp. 36-38, 2008.
- Paulsen, H. and Riegger, U. (2006) SensorGIS - Geodaten in Echtzeit. In: GIS-Business 08/2006, pp. 17-19, Cologne.

- Portele, C. (ed.) (2007) Geography Markup Language (GML).
<http://www.opengeospatial.org>, OpenGIS Standard OGC 07-036, Version 3.2.1,
27 August 2004. (12 September 2009)
- Questex Media Group, Inc. (2009) IBM and Salerno to Develop a Smarter City.
<http://www.sensormag.com/sensors>, Sensors Magazine, 21 May 2009. (4
October 2009)
- Rahm, E., Thor, A. and Aumueller D. (2007) Dynamic Fusion of Web Data. XSym
2007, pp. 14-16, Vienna, Austria.
- Resch, B., Calabrese, F., Ratti, C. and Biderman, A. (2008) An Approach Towards a
Real-time Data Exchange Platform System Architecture. In: Proceedings of the
6th Annual IEEE International Conference on Pervasive Computing and
Communications, Hong Kong, 17-21 March 2008.
- Resch, B., Mittlboeck, M., Girardin, F., Britter, R. and Ratti, C. (in press) Live
Geography – Embedded Sensing for Standardised Urban Environmental
Monitoring. International Journal On Advances in Systems and Measurements,
2(2-3), pp. pending.
- Resch, B., Mittlboeck, M., Girardin, F., Britter, R. and Ratti, C. (2009a) Real-time
Geo-awareness – Sensor Data Integration for Environmental Monitoring in the
City. In: Proceedings of the IEEE International Conference on Advanced
Geographic Information Systems & Web Services – GEOWS2009, pp. 92-97, 1-7
February 2009, Cancun, Mexico.
- Resch, B., Mittlboeck, M., Lipson, S., Welsh, M., Bers, J., Britter, R. and Ratti, C.
(2009b) Urban Sensing Revisited – Common Scents: Towards Standardised
Geo-sensor Networks for Public Health Monitoring in the City. In: Proceedings of
the 11th International Conference on Computers in Urban Planning and Urban
Management – CUPUM2009, 16-18 June 2009, Hong Kong.
- Ricadela, A. (2005) Sensors Everywhere. Information Week, No. 1023, p. 32, 24
January 2005.

- Ride The Lobster (2007) Britech to Track and Map Unicycle Teams as They Ride The Lobster. <http://ridethelobster.blogspot.com>, 8 September 2007. (21 August 2009)
- Rittman, M. (2008) An Introduction to Real-Time Data Integration. <http://www.oracle.com/technology/pub/articles/rittman-odi.html>, 2008. (22 May 2009)
- Riva, O. and Borcea, C. (2007) The Urbanet Revolution: Sensor Power to the People!. IEEE Pervasive Computing, 6(2), pp. 41-49, Apr-Jun 2007.
- Sarjakoski, T., Sester, M., Illert, A., Rystedt, B., Nissen, F. and Ruotsalainen, R. (2004) Geospatial Info-mobility Service by Real-time Data-integration and Generalisation. <http://gimodig.fgi.fi>, 8 November 2004. (22 May 2009)
- Schut, P. (ed.) (2007) Web Processing Service. <http://www.opengeospatial.org>, OpenGIS Standard, Version 1.0.0, OGC 05-007r7, 8 June 2007. (19 June 2009)
- SENSEable City Laboratory (2009) MIT SENSEable City Lab. <http://senseable.mit.edu>, August 2009. (21 August 2009)
- Shi, E. and Perrig, A. (2004) Designing Secure Sensor Networks. IEEE Wireless Communications, 11(6), pp. 38-43, Carnegie Mellon University, Pittsburgh, PA, USA, December 2004.
- Simonis, I. and Wytzisk, A. (2003) Web Notification Service. <http://www.opengeospatial.org>, OpenGIS Consortium Discussion Paper OGC 03-008r2, Version 0.1.0, 21 April 2003. (12 May 2005)
- Simonis, I. (Ed.) (2007a) Sensor Alert Service. <http://www.opengeospatial.org>, Candidate OpenGIS Interface Standard OGC 06-028r5, Version 0.9.0, 14 May 2007. (19 August 2009)
- Simonis, I. (Ed.) (2007b) Sensor Planning Service. <http://www.opengeospatial.org>, OpenGIS Implementation Specification OGC 07-014r3, Version 1.0, 2 February 2007. (18 August 2009)

- Smith, S. (2002) Buzzword for the Millennium. www.gisvisionmag.com, January 2002. (19 April 2005)
- SourceForge (2009) PLWPS Web Processing Service. <http://plwps.sourceforge.net>, 2009. (25 September 2009)
- Srivastava, M., Hansen, M., Burke, J., Parker, A., Reddy, S., Saurabh, G., Allman, M., Paxson, V. and Estrin D. (2006) Wireless Urban Sensing Systems. Technical Report 65, Center for Embedded Network Sensing, UCLA, April 2006.
- Stasch, C. (2009) 52°North SOS Data Models. <http://www.52north.org>, 7 January 2009. (5 September 2009)
- Steel, D. (2005) Smart Dust. www.uhisrc.com, UH Information Systems Research Center - Technology Briefing, March 2005. (28 April 2005)
- Steinberg, A.N. and Bowman, C.L. (2004) Rethinking the JDL Data Fusion Levels. National Symposium on Sensor and Data Fusion, John Hopkins Applied Physics Laboratory, 2004.
- Sybase Inc. (2008) Real-Time Events Data Integration Software. <http://www.sybase.com/products/dataintegration/realtimeevents>, 2008. (22 December 2008)
- Tada, N., Izawa, Y., Kimoto, M., Maruyama, T., Ohno, T. and Nakayama M. (2000) IAA System ("I Am Alive"): The Experiences of the Internet Disaster Drills. <http://www.isoc.org>, Proceedings of INET2000, July 2000. (27 April 2005)
- Takahashi, D. (2004) Smart Dust – Tiny Motion-sensor Technology Offers a World of Opportunities. <http://www.bradenton.com>, 24 October 2004. (2 May 2005)
- Tao, V. (2003) GEO World – Sept 2003 – The Smart Sensor Web – A Revolutionary Leap in Earth Observation. <http://www.geoplace.com>, September 2003. (20 September 2008)
- Tygar, J.D. (2003) The Problem with Privacy. In: Proceedings of the 3rd IEEE Workshop on Internet Applications (WIAPP), California University, Berkeley, CA, USA, ISBN 0-7695-1972-5, pp. 2-8, June 2003.

- UC Berkeley (2009) TinyOS Community Forum || An Open-source OS for the Networked Sensor Regime. <http://www.tinyos.net>, 7 July 2009. (28 August 2005)
- United Nations Population Fund (2007) State of World Population 2007: Unleashing the Potential of Urban Growth. United Nations Population Fund Report, ISBN 978-0897148078, New York, NY, UNFPA, 2007.
- University of Oklahoma (2009) OKCnet. <http://okc.mesonet.org>, March 2009. (12 May 2009)
- US Environmental Protection Agency (2009) Global Earth Observation System of Systems (GEOSS) and the Group on Earth Observations (GEO). <http://www.epa.gov/geoss>, 5 May 2009. (28 July 2009)
- Vivid Solutions (2009) JTS Topology Suite. <http://www.vividsolutions.com>, 2009. (4 October 2009)
- Vretanos, P. (Ed.) (2005) Web Feature Service Implementation Specification. <http://www.opengeospatial.org>, OpenGIS Implementation Specification OGC 04-094, Version 1.1.0, 3 May 2005. (8 September 2009)
- Waltz, E. and Llinas, J. (1990) Multisensor Data Fusion. 488 pp., ISBN 978-0890062777, Artech House Publishers, Boston, MA, USA, 1990.
- Warneke, B., Last, M., Liebowitz, B. and Pister, K.S.J. (2001) Smart Dust: Communicating with a Cubic-Millimeter Computer. IEEE Computer Magazine, vol. 34, pp. 44-51, January 2001.
- Whiteside, A. and Evans, J.D. (Eds.) (2008) Web Coverage Service (WCS) Implementation Standard. <http://www.opengeospatial.org>, OGC Implementation Standard OGC 07-067r5, Version 1.1.2, Version 1.1 Corrigendum 2 Release, 19 March 2008. (3 September 2009)
- WIDE Project (2005) WIDE PROJECT Home Page. <http://www.wide.ad.jp>, April 2005. (27 April 2005)

Xu, N. (2004) A Survey of Sensor Network Applications.

<http://courses.cs.tamu.edu>, Computer Science Department, University of Southern California, 2004. (10 August 2009)

Yoo S., Kim B., Park H., Choi J. and Chun, J. (2003) Realization of Real-time Clinical Data Integration Using Advanced Database Technology. In: Proceedings of AMIA 2003 Annual Symposium, pp. 738-742.

Zardini, M. (Ed.) (2006) Sense of the City: An Alternate Approach to Urbanism. 352 pp., ISBN 3-03778-060-6, Lars Müller Publishers, Baden, Switzerland, 2006.

Ziegler, P. and Dittrich, K. (2004) Three Decades of Data Integration – All Problems Solved?. 1st International IFIP Conference on Semantics of a Networked World, University of Zurich, 2004.

List of Abbreviations

ADT	...	Abstract Data Type
AES	...	Advanced Encryption Standard
AI	...	Artificial Intelligence
API	...	Application Programming Interface
BI	...	Business Intelligence
BP	...	British Petroleum
BPEL	...	Business Process Execution Language
CDMA	...	Code Division Multiple Access
CGI	...	Common Gateway Interface
COP	...	Common Operation Picture
CPM	...	Corporate Performance Management
CQL	...	Continuous Query Language
CRS	...	Coordinate Reference System
CSW	...	Catalogue Service for Web
DARPA	...	Defense Advanced Research Projects
DoS	...	Denial of Service
DSN	...	Distributed Sensor Networks
EAS	...	Emergency Alert System
ebRIM	...	e-business Registry Information Model
EC	...	European Commission
EO	...	Earth Observation
EQL	...	Event Query Language
ESA	...	European Space Agency
FTP	...	File Transfer Protocol
GENESIS	...	GENeric European Sustainable Information Space for environment
GEO	...	Geostationary Orbit
GEO	...	Group of Earth Observations
GEOSS	...	Global Earth Observation System of Systems
GIF	...	Graphics Interchange Format
GIS	...	Geographic Information System
GMES	...	Global Monitoring for Environment and Security
GML	...	Geography Markup Language
GNSS	...	Global Navigation Satellite System
GPS	...	Global Positioning System
GSM	...	Global System for Mobile Communications
HCI	...	Human Computer Interaction

HTTP	...	Hypertext Transfer Protocol
I ² C	...	Inter-Integrated Circuit
IAA	...	I Am Alive
IDW	...	Inverse Distance Weighting
IEEE	...	Institute of Electrical and Electronics Engineers
IETF	...	Internet Engineering Task Force
IHI	...	Integrated Health Index
INSPIRE	...	Infrastructure for Spatial Information in Europe
JDL	...	Joint Directors of Laboratories
JPEG	...	Joint Photographic Experts Group
KML	...	Keyhole Markup Language
LAN	...	Local Area Network
LBS	...	Location-based Service
LEO	...	Low Earth Orbit
LGPL	...	GNU Lesser General Public License
Li-Po	...	Lithium-ion Polymer
LS	...	Location Server
M2M	...	Machine-to-Machine
MAC	...	Medium Access Control
MUC	...	Multi User Chat
NASA	...	National Aeronautics and Space Administration
NCW	...	Network Centric Warfare
NFC	...	Near Field Communication
O&M	...	Observations and Measurements
OGC	...	Opengeospatial Consortium
PHP	...	Hypertext Preprocessor
PNG	...	Portable Network Graphics
REST	...	Representational State Transfer
RFID	...	Radio Frequency Identification
RSS	...	Really Simple Syndication
SAS	...	Sensor Alert Service
SDI	...	Spatial Data Infrastructures
SEIS	...	Shared Environmental Information System
SensorML	...	Sensor Model Language
SFS	...	Simple Features Specification
SMS	...	Short Message Service
SOA	...	Service Oriented Architecture
SOAP	...	Simple Object Access Protocol

SOS	...	Sensor Observation Service
SOSUS	...	Sound Surveillance System
SPS	...	Sensor Planning Service
SQL	...	Structured Query Language
SSL	...	Secure Socket Layer
SVG	...	Scalable Vector Graphics
SWE	...	Sensor Web Enablement
UART	...	Universal Asynchronous Receiver/Transmitter
UDDI	...	Universal Description, Discovery and Integration
UFP	...	Ultrafine Particulate
UMTS	...	Universal Mobile Telecommunications System
UNSDI	...	United Nations Spatial Data Infrastructure
URN	...	Uniform Resource Name
USB	...	Universal Serial Bus
USGS	...	U.S. Geological Survey
XML	...	Extensible Markup Language
W3C	...	World Wide Web Consortium
WCPS	...	Web Coverage Processing Service
WIDE	...	Widely Integrated Distributed Environment
WiFi	...	Wireless Fidelity
WNS	...	Web Notification Service
WPS	...	Web Processing Service
WSDL	...	Web Service Description Language
WSS	...	Web Services Security

List of Figures

Figure 1: Urban Monitoring – the Connection between the Physical and the Digital World. (SENSEable City Lab, 2009).....	3
Figure 2: Current Approach – Monolithic Single-purpose Sensor Systems.	5
Figure 3: Feedback Loop for Environmental Monitoring in the City.	23
Figure 4: Basic Sensor Web Concept. (Delin et al., 2005)	25
Figure 5: Sensor Web with Inter-communicating Sensors. (Liang et al., 2004)	26
Figure 6: Standardisation Bodies Classified by Legal Bindingness and Genericness.	27
Figure 7: “Spec” Sensor Pod next to the Tip of a Ball Point Pen. (Steel, 2005)	35
Figure 8: Decision Tree of Sensor Web Architectures.	55
Figure 9: Basic Architecture Components of the Live Geography Infrastructure....	66
Figure 10: Live Geography – Workflow Illustration with Interfaces of the Architecture Approach.	67
Figure 11: Web Service Standards Structure. (IBM Corporation, 2005).....	69
Figure 12: General SWE Idea and Architecture. (Botts et al., 2007)	70
Figure 13: Element Hierarchy in SensorML. (Botts, 2004).....	71
Figure 14: General SOS Usage Concept. (Adapted from Na and Priest, 2007)	72
Figure 15: General SAS Functionality. (Simonis, 2007a).....	73
Figure 16: Functional Connection of SWE Components. (Adapted from Botts et al., 2007).....	76
Figure 17: WFS Protocol Diagram. (Computational Geoscience Group, 2009)	78
Figure 18: Class Hierarchy of GML 3.1.0. (International Organization for Standardization, 2005a).....	80
Figure 19: JDL Data Fusion Levels. (Steinberg and Bowman, 2004)	85
Figure 20: ESRI Geo-processing Scheme. (ESRI, 2009)	87
Figure 21: Sensor Fusion and Data Fusion/Geo-processing in the Geo-Workflow Chain.	88
Figure 22: Chaining of Several Sensor Fusion and Data Fusion/Geo-processing Operations.	89
Figure 23: General CEP Architecture with Data Processing Levels. (Adapted from Bass, 2007)	93
Figure 24: Communication of the Sensor Device with other Components in the System.....	102
Figure 25: Gumstix Verdex Pro XM4 with GPS and a Combined Air Temperature and Humidity Sensor.	104
Figure 26: Software Infrastructure on the Embedded Sensing Device.	106

Figure 27: 52°North Relation Database Schema for Mobile Sensors. (Stasch, 2009)	109
Figure 28: Relational Schema of the Embedded Measurement Database Structure.	110
Figure 29: OGC Data Model for O&M.	112
Figure 30: Sequence Diagram of a Feature Query within the Custom Data Store.	116
Figure 31: Internal Service Stack of the CEP Component.	121
Figure 32: 2D Interpolation Using the Inverse Distance Weighting (IDW) Algorithm.	124
Figure 33: Spatio-temporal CO ₂ Data Analysis Using Tracking Analyst.	125
Figure 34: Web-based Interpolation of Real-time Sensor Data.	126
Figure 35: Live Geography: Standards-based Loosely-coupled Service Infrastructure.	132
Figure 36: Time Series Visualisation of Pollutant Measurements.	139
Figure 37: Interpolation of Temperature Values for the Detection of Urban Heat Islands.	140
Figure 38: Biometric Parameter History with Geographical Location Illustration.	141
Figure 39: Web-based Live Geo-processing for Environmental Monitoring.	143
Figure 40: Workflow and Processing Approaches for Sensor Fusion Data.	149
Figure 41: Fully-fledged Live Geography Infrastructure.	150

List of Tables

Table 1: Definitions for Data Fusion and Geo-processing	84
Table 2: Segments in the Sequence Diagram for the GeoServer Data Store Sensor Query.....	115
Table 3: CEP Frameworks.	119
Table 4: Geographical Code Libraries.....	120

Annex I – XML Structures

OGC SOS Version 1.0.0

SOS 1.0.0 GetCapabilities Request

GET-Request

http://iSPACEEnvironment005.researchstudio.at/cgi-bin/sos?
request=GetCapabilities&
version=1.0.0&
service=SOS

POST-Request

```
<?xml version="1.0" encoding="UTF-8"?>
<GetCapabilities xmlns="http://schemas.opengis.net/sos/1.0.0"
  xmlns:ows="http://schemas.opengis.net/ows/1.1.0"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://schemas.opengis.net/sos/1.0.0
http://schemas.opengis.net/sos/1.0.0/sosGetCapabilities.xsd" service="SOS"
  updateSequence="2009-06-01T11:12:39+02">
  <ows:AcceptVersions>
    <ows:Version>1.0.0</ows:Version>
  </ows:AcceptVersions>
  <ows:Sections>
    <ows:Section>OperationsMetadata</ows:Section>
    <ows:Section>ServiceIdentification</ows:Section>
    <ows:Section>Filter_Capabilities</ows:Section>
    <ows:Section>Contents</ows:Section>
  </ows:Sections>
</GetCapabilities>
```

SOS 1.0.0 GetCapabilities Response

```
<?xml version="1.0" encoding="utf-8"?>
<Capabilities xmlns="http://schemas.opengis.net"
  xmlns:gml="http://schemas.opengis.net/gml/3.1.1"
  xmlns:swe="http://schemas.opengis.net/sweCommon/1.0.1"
  xmlns:sos="http://schemas.opengis.net/sos/1.0.0"
  xmlns:om="http://schemas.opengis.net/om/1.0.0"
  xmlns:tml="http://schemas.opengis.net/tml/1.0.0"
  xmlns:sml="http://schemas.opengis.net/sensorML/1.0.1"
  xmlns:ows="http://schemas.opengis.net/ows/1.1.0/"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:xlink="http://www.w3.org/1999/xlink">
```

```

xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns:ogc="http://schemas.opengis.net"
xsi:schemaLocation="http://schemas.opengis.net/sweCommon/1.0.1
http://schemas.opengis.net/sweCommon/1.0.1/swe.xsd"
updateSequence="2009-06-01T11:12:39+02" version="1.0.0">
  <ows:ServiceIdentification>
    <ows:Title>Live Geography SOS</ows:Title>
    <ows:Abstract>Live Geography Sensor Observation Server</ows:Abstract>
    <ows:Keywords>
      <ows:keyword>Live Geography, SOS</ows:keyword>
    </ows:Keywords>
    <ows:ServiceType
      codeSpace="http://opengeospatial.net">OGC:SOS</ows:ServiceType>
    <ows:ServiceTypeVersion>1.0.0</ows:ServiceTypeVersion>
    <ows:Fees>NONE</ows:Fees>
    <ows:AccessConstraints>NONE</ows:AccessConstraints>
  </ows:ServiceIdentification>
  <ows:ServiceProvider>
    <ows:ProviderName>Bernd Resch</ows:ProviderName>
    <ows:ProviderSite xlink:href="http://www.berndresch.com"/>
    <ows:ServiceContact>
      <ows:IndividualName>Bernd Resch</ows:IndividualName>
      <ows:PositionName>Research Scientist, PhD
        Student</ows:PositionName>
      <ows:ContactInfo>
        <ows:Phone>
          <ows:Voice>+43662908585220</ows:Voice>
          <ows:Facsimile>+43662908585222</ows:Facsimile>
        </ows:Phone>
        <ows:Address>
          <ows:DeliveryPoint>Bernd Resch</ows:DeliveryPoint>
          <ows:City>Salzburg</ows:City>
          <ows:AdministrativeArea>Salzburg</ows:AdministrativeArea>
          <ows:PostalCode>5020</ows:PostalCode>
          <ows:Country>Austria</ows:Country>
          <ows:ElectronicMailAddress>
            bernd.resch@researchstudio.at</ows:ElectronicMailAddress>
          </ows:Address>
        </ows:ContactInfo>
      </ows:ServiceContact>
    </ows:ServiceProvider>
    <ows:OperationsMetadata>
      <ows:Operation name="GetCapabilities">
        <ows:DCP>
          <ows:HTTP>
            <ows:Get xlink:href="https://
              iSPACEEnvironment005.researchstudio.at/cgi-bin/sos"/>
            <ows:Post xlink:href="https://
              iSPACEEnvironment005.researchstudio.at/cgi-bin/sos"/>
          </ows:HTTP>
        </ows:DCP>
        <ows:Parameter name="request" use="required">
          <ows:Value>GetObservation</ows:Value>
        </ows:Parameter>
        <ows:Parameter name="version" use="required">
          <ows:Value>1.0.0</ows:Value>
          <ows:Value>0.3.4</ows:Value>
        </ows:Parameter>
      </ows:Operation>
    </ows:OperationsMetadata>
  </ows:ServiceMetadata>
</ows:Service>

```

```

    </ows:Parameter>
  </ows:Operation>
  <ows:Operation name="GetObservation">
    <ows:DCP>
      <ows:HTTP>
        <ows:Get xlink:href="https://
          iSPACEEnvironment005.researchstudio.at/cgi-bin/sos"/>
        <ows:Post xlink:href="https://
          iSPACEEnvironment005.researchstudio.at/cgi-bin/sos"/>
      </ows:HTTP>
    </ows:DCP>
  </ows:Operation>
  <ows:Operation name="DescribeSensor">
    <ows:DCP>
      <ows:HTTP>
        <ows:Get xlink:href="https://
          iSPACEEnvironment005.researchstudio.at/cgi-bin/sos"/>
        <ows:Post xlink:href="https://
          iSPACEEnvironment005.researchstudio.at/cgi-bin/sos"/>
      </ows:HTTP>
    </ows:DCP>
    <ows:Parameter name="procedure" use="required">
      <om:procedure xlink:href="
        urn:ogc:object:Sensor:iSPACE:AirTemperature"/>
      <om:procedure xlink:href="
        urn:ogc:object:Sensor:iSPACE:RelativeHumidity"/>
    </ows:Parameter>
  </ows:Operation>
</ows:OperationsMetadata>
<ogc:Filter_Capabilities>
  <ogc:Temporal_Capabilities>
    <ogc:TemporalOperands>
      <ogc:TemporalOperand>gml:TimeInstant</ogc:TemporalOperand>
      <ogc:TemporalOperand>gml:TimePeriod</ogc:TemporalOperand>
    </ogc:TemporalOperands>
  </ogc:Temporal_Capabilities>
  <ogc:Scalar_Capabilities>
    <ogc:ComparisonOperators>
      <ogc:ComparisonOperator>LessThan
    </ogc:ComparisonOperator>
      <ogc:ComparisonOperator>GreaterThan
    </ogc:ComparisonOperator>
      <ogc:ComparisonOperator>LessThanEqualTo
    </ogc:ComparisonOperator>
      <ogc:ComparisonOperator>GreaterThanEqualTo
    </ogc:ComparisonOperator>
      <ogc:ComparisonOperator>EqualTo
    </ogc:ComparisonOperator>
      <ogc:ComparisonOperator>NotEqualTo
    </ogc:ComparisonOperator>
      <ogc:ComparisonOperator>Between
    </ogc:ComparisonOperator>
    </ogc:ComparisonOperators>
  </ogc:Scalar_Capabilities>
</ogc:Filter_Capabilities>
<Contents>
  <ObservationOfferingList>

```

```

<ObservationOffering gml:id="environmental_parameters">
  <gml:description>Measuring different environmental
    parameters</gml:description>
  <gml:name>
    urn:iSPACE:offering:environmental_parameters
  </gml:name>
  <gml:srsName>urn:ogc:def:crs:OGC:1.3:EPSG:6.7:4326
  </gml:srsName>
  <gml:boundedBy>
    <gml:Envelope
      srsName="urn:ogc:def:crs:OGC:1.3:EPSG:6.7:4326">
      <gml:lowerCorner>13.038785 13.038785</gml:lowerCorner>
      <gml:upperCorner>13.038785 13.038785</gml:upperCorner>
    </gml:Envelope>
  </gml:boundedBy>
  <time>
    <gml:TimePeriod>
      <gml:beginPosition>1999-12-31T23:00:00</gml:beginPosition>
      <gml:endPosition>2009-12-31T23:00:00</gml:endPosition>
    </gml:TimePeriod>
  </time>
  <om:procedure
    xlink:href="urn:ogc:object:Sensor:iSPACE:AirTemperature"/>
  <om:procedure
    xlink:href="urn:ogc:object:Sensor:iSPACE:RelativeHumidity"/>
  <om:observedProperty xlink:href=
    "urn:x-ogc:def:phenomenon:OGC:AirTemperature"/>
  <om:observedProperty xlink:href=
    "urn:x-ogc:def:phenomenon:OGC:RelativeHumidity"/>
  <om:featureOfInterest
    xlink:href="urn:ogc:def:feature:iSPACE:FOI_Environment"/>
  <responseFormat>text/xml; subtype="om/1.0.0"</responseFormat>
  <resultModel>om:Observation</resultModel>
  <responseMode>inline</responseMode>
</ObservationOffering>
</ObservationOfferingList>
</Contents>
</Capabilities>

```

SOS 1.0.0 DescribeSensor Request

```
<?xml version="1.0" encoding="UTF-8"?>
<DescribeSensor xmlns="http://schemas.opengis.net/sas/0.9"
  xmlns:sml="http://schemas.opengis.net/sensorML/1.0.1"
  xmlns:xsl="http://www.w3.org/2001/XMLSchema-instance" service="SOS"
  version="1.0.0">
  <SensorID>urn:x-ogc:object:sensor:iSPACE:Environment:005</SensorID>
</DescribeSensor>
```

SOS 1.0.0 DescribeSensor Response

```
<?xml version="1.0" encoding="utf-8"?>
<sml:Sensor xmlns:swe="http://schemas.opengis.net/sweCommon/1.0.1"
  xmlns:xlink="http://www.w3.org/1999/xlink" version="1.0"
  xmlns:sml="http://schemas.opengis.net/sensorML/1.0.1"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  id="urn:ogc:object:feature:Sensor:iSPACE:RelativeHumidity
urn:ogc:object:feature:Sensor:iSPACE:AirTemperature ">
  <sml:identification>
    <sml:IdentifierList>
      <sml:identifier name="longName">
        <sml:Term definition=
          "urn:ogc:object:feature:Sensor:iSPACE:AirTemperature "/>
      </sml:identifier>
    </sml:IdentifierList>
  </sml:identification>
  <sml:inputs>
    <sml:InputList>
      <sml:input name=" RelativeHumidity ">
        <swe:Quantity definition=
          "urn:x-ogc:def:phenomenon:OGC:AirTemperature "
          uom="urn:ogc:unit:degC"/>
      </sml:input>
    </sml:InputList>
  </sml:inputs>
  <sml:outputs>
    <sml:OutputList>
      <sml:input name="AirTemperature ">
        <swe:Quantity definition=
          "urn:x-ogc:def:phenomenon:OGC:AirTemperature "
          uom="urn:ogc:unit:degC"/>
      </sml:input>
    </sml:OutputList>
  </sml:outputs>
  <sml:identification>
    <sml:IdentifierList>
      <sml:identifier name="longName">
        <sml:Term definition=
          "urn:ogc:object:feature:Sensor:iSPACE:RelativeHumidity "/>
      </sml:identifier>
    </sml:IdentifierList>
  </sml:identification>
  <sml:inputs>
    <sml:InputList>
```

```

    <sml:input name=" RelativeHumidity ">
      <swe:Quantity definition=
        "urn:x-ogc:def:phenomenon:OGC:RelativeHumidity"
        uom="urn:ogc:unit:degC"/>
    </sml:input>
  </sml:InputList>
</sml:inputs>
<sml:outputs>
  <sml:OutputList>
    <sml:input name="RelativeHumidity">
      <swe:Quantity definition=
        "urn:x-ogc:def:phenomenon:OGC:RelativeHumidity"
        uom="urn:ogc:unit:percent"/>
    </sml:input>
  </sml:OutputList>
</sml:outputs>
<sml:referenceFrame>
  <sml:LocalTimeCRS id="SENSOR_TIME">
    <sml:srsName>GMT</sml:srsName>
    <sml:usesCS xlink:href="urn:ogc:crs:GMT"/>
  </sml:LocalTimeCRS>
</sml:referenceFrame>
</sml:Sensor>

```

SOS 1.0.0. GetObservation Request

GET-Request

```

http://iSPACEEnvironment005.researchstudio.at/cgi-bin/sos?
request=GetObservation&
version=1.0.0&
offering=environmental_parameters&
observedProperty=AirTemperature,RelativeHumidity&
service=SOS

```

POST-Request

```

<?xml version="1.0" encoding="UTF-8"?>
<GetObservation xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:ogc="http://schemas.opengis.net"
  xmlns:sos="http://schemas.opengis.net/sos/1.0.0"
  xmlns:om="http://schemas.opengis.net/om/1.0.0"
  xsi:schemaLocation="http://schemas.opengis.net/sweCommon/1.0.1
http://schemas.opengis.net/sweCommon/1.0.1/swe.xsd" service="SOS"
  version="1.0.0" srsName="EPSG:4326">
  <sos:offering>urn:iSPACE:offering:environmental_parameters</sos:offering>
  <observedProperty>AirTemperature</observedProperty>
  <observedProperty>RelativeHumidity</observedProperty>
</GetObservation>

```

SOS 1.0.0 GetObservation Response

```
<?xml version="1.0" encoding="utf-8"?>
<om:Observation xmlns="http://schemas.opengis.net/om/1.0.0"
  xmlns:gml="http://schemas.opengis.net/gml/3.1.1"
  xmlns:swe="http://schemas.opengis.net/sweCommon/1.0.1"
  xmlns:sos="http://schemas.opengis.net/sos/1.0.0"
  xmlns:ows="http://schemas.opengis.net/ows/1.1.0/"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:xlink="http://www.w3.org/1999/xlink"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:ogc="http://schemas.opengis.net"
  xsi:schemaLocation="http://schemas.opengis.net/sweCommon/1.0.1
http://schemas.opengis.net/sweCommon/1.0.1/swe.xsd">
  <gml:description>Environmental Parameter Measurement</gml:description>
  <gml:name>urn:iSPACE:offering:environmental_parameters </gml:name>
  <om:samplingTime>
    <gml:TimePeriod>
      <gml:beginPosition>1999-12-31T23:00:00</gml:beginPosition>
      <gml:endPosition>2009-12-31T23:00:00</gml:endPosition>
    </gml:TimePeriod>
  </om:samplingTime>
  <om:procedure xlink:href="urn:ogc:object:Sensor:iSPACE:AirTemperature"/>
  <om:procedure xlink:href="urn:ogc:object:Sensor:iSPACE:RelativeHumidity"/>
  <om:observedProperty>
    <swe:DataRecord>
      <swe:field name="time">
        <swe:Time definition="urn:ogc:def:phenomenon:time:iso8601"/>
      </swe:field>
      <swe:field name="longitude">
        <swe:Quantity
          definition="urn:ogc:phenomenon:location:EPSG:4326:longitude">
          <swe:uom xlink:href="urn:ogc:def:uom:OGC:deg"/>
        </swe:Quantity>
      </swe:field>
      <swe:field name="latitude">
        <swe:Quantity
          definition="urn:ogc:phenomenon:location:EPSG:4326:latitude">
          <swe:uom xlink:href="urn:ogc:def:uom:OGC:deg"/>
        </swe:Quantity>
      </swe:field>
      <swe:field name="AirTemperature">
        <swe:Quantity definition="urn:x-
ogc:def:phenomenon:OGC:AirTemperature">
          <swe:uom xlink:href="urn:ogc:def:uom:OGC:degC"/>
        </swe:Quantity>
      </swe:field>
      <swe:field name="RelativeHumidity">
        <swe:Quantity definition="urn:x-
ogc:def:phenomenon:OGC:RelativeHumidity">
          <swe:uom xlink:href="urn:ogc:def:uom:OGC:percent"/>
        </swe:Quantity>
      </swe:field>
    </swe:DataRecord>
  </om:observedProperty>
</om:Observation>
```

```

</om:observedProperty>
<om:featureOfInterest
  xlink:href="urn:ogc:def:feature:iSPACE:FOI_Environment"/>
<om:result>
  <swe:DataArray>
    <swe:elementCount>
      <swe:Count>
        <swe:value>10</swe:value>
      </swe:Count>
    </swe:elementCount>
    <swe:elementType>
      <swe:CompositePhenomenon>
        <swe:component xlink:href="urn:x-
          ogc:def:phenomenon:OGC:AirTemperature"/>
        <swe:component xlink:href="urn:x-
          ogc:def:phenomenon:OGC:RelativeHumidity"/>
      </swe:CompositePhenomenon>
    </swe:elementType>
    <swe:encoding>
      <swe:TextBlock decimalSeparator="." tokenSeparator=","
        blockSeparator="@@"/>
    </swe:encoding>
    <swe:values>
      2009-09-05T15:57:34Z,13.038785,13.038785,
      27.760000,28.976608@@
      2009-09-05T15:57:39Z,13.038785,13.038785,
      27.760000,28.976608@@
      2009-09-05T15:57:44Z,13.038785,13.038785,
      27.760000,28.904787@@
      2009-09-05T15:57:50Z,13.038785,13.038785,
      27.760000,28.904787@@
      2009-09-05T15:57:55Z,13.038785,13.038785,
      27.760000,28.976608@@
      2009-09-05T15:58:00Z,13.038785,13.038785,
      27.750000,28.975821@@
      2009-09-05T15:58:05Z,13.038785,13.038785,
      27.750000,28.975821@@
      2009-09-05T15:58:10Z,13.038785,13.038785,
      27.750000,28.975821@@
      2009-09-05T15:58:15Z,13.038785,13.038785,
      27.760000,28.976608@@
      2009-09-05T15:58:20Z,13.038785,13.038785,
      27.750000,28.975821
    </swe:values>
  </swe:DataArray>
</om:result>
</om:Observation>

```

OGC SAS Version 0.9

SAS 0.9 DescribeAlert Request

```
<?xml version="1.0" encoding="UTF-8"?>
<DescribeAlert xmlns="http://schemas.opengis.net/sas/0.9"
  xmlns:sml="http://schemas.opengis.net/sensorML/1.0.1"
  xmlns:xsl="http://www.w3.org/2001/XMLSchema-instance" service="SAS"
  version="0.9">
  <SensorID>urn:x-ogc:object:sensor:iSPACE:Environment:005</SensorID>
</DescribeAlert>
```

SAS 0.9 DescribeAlert Response

```
<?xml version="1.0" encoding="utf-8"?>
<DescribeAlertResponse>
  <messageStructure>
    <swe:DataBlockDefinition>
      <swe:components name="sensor data structure">
        <swe:DataRecord
          xsi:schemaLocation="http://schemas.opengis.net/sweCommon/1.0.1
            http://schemas.opengis.net/sweCommon/1.0.1/swe.xsd">
          <swe:field name="time">
            <swe:Time
              definition="urn:ogc:def:phenomenon:time:iso8601"/>
          </swe:field>
          <swe:field name="longitude">
            <swe:Quantity definition=
              "urn:ogc:phenomenon:location:EPSG:4326:longitude">
              <swe:uom xlink:href="urn:ogc:def:uom:OGC:deg"/>
            </swe:Quantity>
          </swe:field>
          <swe:field name="latitude">
            <swe:Quantity definition=
              "urn:ogc:phenomenon:location:EPSG:4326:latitude">
              <swe:uom xlink:href="urn:ogc:def:uom:OGC:deg"/>
            </swe:Quantity>
          </swe:field>
          <swe:field name="AirTemperature">
            <swe:Quantity definition="urn:x-
              ogc:def:phenomenon:OGC:AirTemperature">
              <swe:uom xlink:href="urn:ogc:def:uom:OGC:degC"/>
            </swe:Quantity>
          </swe:field>
          <swe:field name="RelativeHumidity">
            <swe:Quantity definition="urn:x-
              ogc:def:phenomenon:OGC:RelativeHumidity">
              <swe:uom xlink:href="urn:ogc:def:uom:OGC:percent"/>
            </swe:Quantity>
          </swe:field>
        </swe:DataRecord>
      </swe:components>
    <swe:encoding>
```

```

        <swe:TextBlock decimalSeparator="." blockSeparator="@@"  

            tokenSeparator=" "/>  

    </swe:encoding>  

</swe:DataBlockDefinition>  

</messageStructure>  

</DescribeAlertResponse>

```

SAS 0.9 Subscribe Request

```

<?xml version="1.0" encoding="UTF-8"?>  

<Subscribe xmlns="http://schemas.opengis.net/sas/0.9"  

  xmlns:sml="http://schemas.opengis.net/sensorML/1.0.1"  

  xmlns:xsl="http://www.w3.org/2001/XMLSchema-instance" service="SAS"  

  version="0.9">  

  <EventFilter>  

    <ValueFilterList>  

      <member>  

        <ValueFilter definition=  

          "urn:x-goc:def:phenomenon:OGC:RelativeHumidity">  

          <filterCriteria>  

            <isGreaterThanOrEqualTo>80</isGreaterThanOrEqualTo>  

            <isSmallerThan>100</isSmallerThan>  

          </filterCriteria>  

        </ValueFilter>  

        <ValueFilter definition=  

          "urn:x-goc:def:phenomenon:OGC:AirTemperature">  

          <filterCriteria>  

            <isGreaterThanOrEqualTo>70</isGreaterThanOrEqualTo>  

            <isSmallerThan>110</isSmallerThan>  

          </filterCriteria>  

        </ValueFilter>  

      </member>  

    </ValueFilterList>  

  </EventFilter>  

</Subscribe>

```

SAS 0.9 Alert

```

<?xml version="1.0" encoding="UTF-8"?>  

<Alert xmlns="http://schemas.opengis.net/sas/0.9"  

  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"  

  xsi:schemaLocation="http://schemas.opengis.net/sas/0.9 ../sasCommon.xsd">  

  <SensorID>urn:x-ogc:object:sensor:iSPACE:Environment:005</SensorID>  

  <Timestamp>2009-10-16T17:29:09Z</Timestamp>  

  <AlertData>13.038743 47.823887 22.70 38.50</AlertData>  

</Alert>

```

SAS 0.9 RenewSubscription Request

```
<?xml version="1.0" encoding="UTF-8"?>
<RenewSubscription xmlns="http://schemas.opengis.net/sas/0.9"
  xmlns:sml="http://schemas.opengis.net/sensorML/1.0.1"
  xmlns:xsl="http://www.w3.org/2001/XMLSchema-instance" service="SAS"
  version="0.9">
  <SubscriptionID>env1</SubscriptionID>
</RenewSubscription>
```

SAS 0.9 CancelSubscription Request

```
<?xml version="1.0" encoding="UTF-8"?>
<CancelSubscription xmlns="http://schemas.opengis.net/sas/0.9"
  xmlns:sml="http://schemas.opengis.net/sensorML/1.0.1"
  xmlns:xsl="http://www.w3.org/2001/XMLSchema-instance" service="SAS"
  version="0.9">
  <SubscriptionID>env1</SubscriptionID>
</CancelSubscription>
```