

Urban Sensing Revisited – Common Scents: Towards Standardised Geo-sensor Networks for Public Health Monitoring in the City

Bernd RESCH
Research Scientist ^{1,2}
E-mail: berno [at] mit.edu
Tel: +43-662-908585-220

Manfred MITTLBOECK
Key Researcher ¹

Sam LIPSON
Director ³

Matt WELSH
Associate Professor ⁴

Josh BERS
Senior Software Engineer ⁵

Rex BRITTER
Visiting Professor ²

Carlo RATTI
Director, Associate Professor ²

1 Research Studios Austria – iSPACE
2 MIT – SENSEable City Lab
3 City of Cambridge, MA US – Public Health Department
4 Harvard University – School of Engineering and Applied Sciences
5 BBN Technologies – Advanced Networking

Abstract: Pervasive environmental monitoring implies a wide range of technical, but also socio-political challenges, and this applies especially to the sensitive context of the city. In this paper, we elucidate issues for bringing out pervasive urban sensor networks and associated concerns relating to fine-grained information provision. We present the *Common Scents* project, which is based on the *Real-time Geo-awareness* approach, and show how it can overcome these challenges. As opposed to hitherto sensing networks, which are mostly built up in monolithic and closed systems, the *Common Scents* approach aims to establish an open, standards based and modular infrastructure. This ensures interoperability, portability and flexibility, which are crucial prerequisites for pervasive urban sensing. The implementation – a real-time data integration and analysis system for air quality assessment – has been realised on top of the CitySense sensor network in the City of Cambridge, MA US together with the city's Public Health Department responding to concrete needs of the city and its inhabitants.

Keywords: environmental monitoring, pervasive urban sensing, GIS data analysis, standardised geo-sensor networks, public health

1. INTRODUCTION

Pervasive environmental sensing encompasses sensors, communications, massive data manipulation and analysis, data fusion with mathematical modelling, the production of outputs on a variety of scales from the street or home to the global, the provision of information as both hard data and user-sensitive visualisation, together with appropriate delivery structures. This paper addresses the nature of this supply

chain; one overarching aspect is that all elements 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 chapter 2. 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 environmental regulatory structure moves from a mathematical modelling base to a more pervasive monitoring structure.

However of specific interest in this paper is our concern that most sensor networks are being built up in monolithic and specific application-centred measurement systems. In consequence, there is a clear gap between sensor network research and mostly very heterogeneous end user requirements. Sensor network research is often dedicated to a long-term vision, which tells a compelling story about potential applications. On the contrary, the actual implementation is mostly not more than a very limited demonstration without taking into account well-known issues such as interoperability, sustainable development, portability or the coupling with established data analysis systems. In consequence, the availability of wireless sensor networks is limited. Deborah Estrin pointed out in a talk in 2004 that no real sensor network applications exist, apart from short-lived prototypical and very domain-specific demos; and this situation hasn't changed substantially over the last years. There are some experimental examples (s. chapter 2), but the clear social, health and economic benefits have not been demonstrated or described in a manner that would compel this sort of investment.

The goal of the *Common Scents* project is that its highly flexible architecture will bring sensor network applications one step further towards the realisation of the vision of a "digital skin for planet earth" and have particularly far-reaching impacts on urban monitoring systems through the deployment of ubiquitous and very fine-grained sensor networks. In other words, the broad goal of the project is to develop an overarching infrastructure for various kinds of sensor network applications.

The structure of the paper is as follows: after this introduction and a discussion of related work, we present the *Common Scents* approach and implementation for public health monitoring in chapter 3. Then, we illustrate some possible application areas for the system. In chapters 5 and 6, we treat challenges for the deployment of sensor networks that are specific to the city and analyse how such systems can change the city as the functional context by adding new unseen information layers. We conclude by a short summary, discussion and future outlook.

2. RELATED WORK

As this paper treats a wide variety of research areas, related work in different fields has to be investigated.

The first domain is **sensor network development for environmental 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 very 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.

Paulsen (2008) presents a sensing infrastructure called *PermaGIS* that attempts to combine sensor systems and GIS-based visualisation technologies. The sensing devices, which measure rock temperature at ten minute intervals, 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. The visualisation component uses a number of open standards (OGC WMS, WFS) and open-source services (UMN Map Server, Mapbender).

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 GIS analysis, the system architecture is not based on open (geospatial) standards, only standard web services. The web portal implementation, called SensorMap (Nath et al., 2006), uses the Sensor Description Markup Language (SDML), an application-specific dialect of the OGC SensorML standard.

A GIS mash-up for environmental data visualisation is presented in the nowCOAST application (National Oceanic and Atmospheric Administration, 2008). Data from providers such as National Oceanic and Atmospheric Administration (NOAA), U.S. Geological Survey (USGS), Department of Defence (DOD) or local airports, 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 with 0.5 degree spatial resolution, produced using a two-dimensional variation interpolation scheme. The system does not make use of open standards for sensor measurements and data provision.

The second related research area is **real-time data integration for GIS** analysis systems. Harrie (2004) and Lehto and Sarjakoski (2005) present web services based on the classic request/response model. Although both methods widely use open GIS standards, they are not suitable for the integration of real-time data for large volumes of data. Sarjakoski et al. (2004) establish a real-time spatial data infrastructure (SDI), which performs several application-specific steps, such as coordinate transformation, spatial data generalisation, query processing or map rendering and adaptation. However, the implemented system accounts neither for event-based push mechanisms nor for the integration of sensor data.

Other approaches for real-time data integration try to tackle the issue from a database perspective. Oracle's system, presented by Rittman (2008), is essentially a middleware between (web) services and a continuously updated database layer. Like Sybase Inc. (2008), the Oracle approach is able to detect database events in order to analyse heterogeneous data sources and trigger actions accordingly. Rahm et al. (2007) present a more dynamic method of data integration and fusion using on-the-fly object matching and metadata repositories to create a flexible data integration environment.

The third related research area is the development of an **open data integration**

system architecture in a non-application specific infrastructure. Srivastava et al. (2006) and Balazinska et al. (2007) present general concepts in a systems architecture and data integration approach but there are no concrete conclusions as to how the final goal of establishing such an infrastructure could be achieved. A more technical approach for ad-hoc sensor networks is described by Riva and Borcea (2007), where the authors discuss challenges to making heterogeneous sensor measurements combinable through the creation of highly flexible middleware components. The method is application-motivated and thus very detailed as far as specific implementation details are concerned. Wood et al. (2007) present a system, which aims.

3. COMMON SCENTS

Within the *Common Scents* project we are looking at developing an interoperable open standards based infrastructure for providing fine-grained air quality data to allow users to assess environmental conditions instantaneously and intuitively. The aim is to provide citizens with unseen up-to-date information layers in order to support their short-term decisions. To achieve this vision, we utilise the CitySense sensor testbed and technologically build our system upon the *Real-time Geo-awareness* approach presented by Resch et al. (2009). It shall be stated that the term “real-time” in our 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.

3.1 Design Principles

In the construction of our technical infrastructure we accounted for different design principles (Service Oriented Architectures – SOA, modular software infrastructures etc.) to ensure flexibility, reusability and portability of the components and the overall infrastructure. Figure 1 shows the modular architecture and the interfaces between the different components.

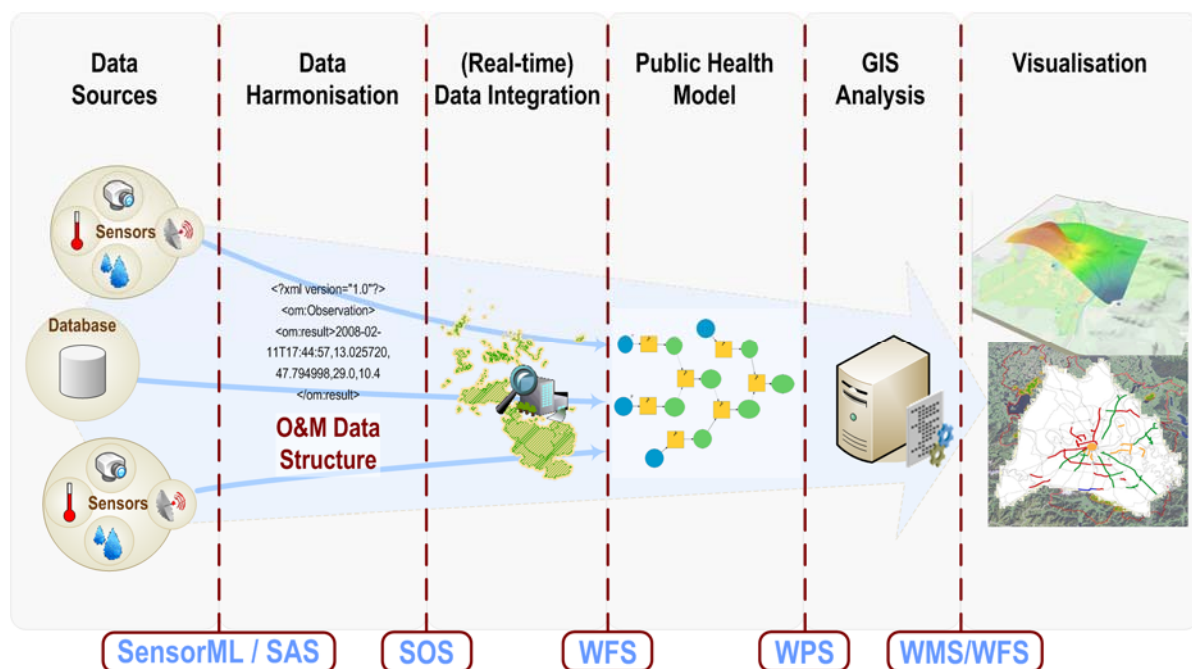


Figure 1 Common Scents – Modular Standardised Infrastructure.

Following principles of SOA and sustainable infrastructure development, we conceived data collection, processing and information provision architecture for the use case of public health, which covers the whole process chain from sensor network development via measurement integration, data analysis and information visualisation, as shown in Figure 1. Thus, our approach builds the architectural bridge between domain-independent sensor network development and use case specific requirements for end user tailored information output.

3.2 Standardised Interfaces

The components of the process chain shown in are separated by several interfaces, which are defined using open standards. The first central group of standards is subsumed under the term Sensor Web Enablement (SWE), an initiative by the OGC that aims to make sensors discoverable, query-able, and controllable over the Internet (Botts et al., 2006). Currently, the SWE family consists of seven standards:

- *Sensor Model Language (SensorML)* – This standard provides an XML schema for defining the geometric, dynamic and observational characteristics of a sensor. Thus, SensorML assists in the discovery of different types of sensors, and supports the processing and analysis of the retrieved data, as well as the geo-location and tasking of sensors.
- *Observations & Measurements (O&M)* – O&M provides a description of sensor observations in the form of general models and XML encodings. This framework labels several terms for the measurements themselves as well as for the relationship between them. Measurement results are expressed as quantities, categories, temporal or geometrical values as well as arrays or composites of these.
- *Transducer Model Language (TML)* – Generally speaking, TML can be understood as O&M's pendant or streaming data by providing a method and message format describing how to interpret raw transducer data.
- *Sensor Observation Service (SOS)* – SOS provides a standardised web service interface allowing access to sensor observations and platform descriptions.
- *Sensor Planning Service (SPS)* – SPS offers an interface for planning an observation query. In effect, the service performs a feasibility check during the set up of a request for data from several sensors.
- *Sensor Alert Service (SAS)* – SAS can be seen as an event-processing engine whose purpose is to identify pre-defined events such as the particularities of sensor measurements, and then generate and send alerts in a standardised protocol format.
- *Web Notification Service (WNS)* – The Web Notification Service is responsible for delivering generated alerts to end-users by E-mail, over HTTP, or via SMS. Moreover, the standard provides an open interface for services, through which a client may exchange asynchronous messages with one or more other services.
- *Sensor Web Registry* – 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 CS-W (Web Catalogue Service) into SWE by building profiles in ebRIM/ebXML (e-business Registry Information Model).

The functional connections between the SWE standards are illustrated in Figure 2.

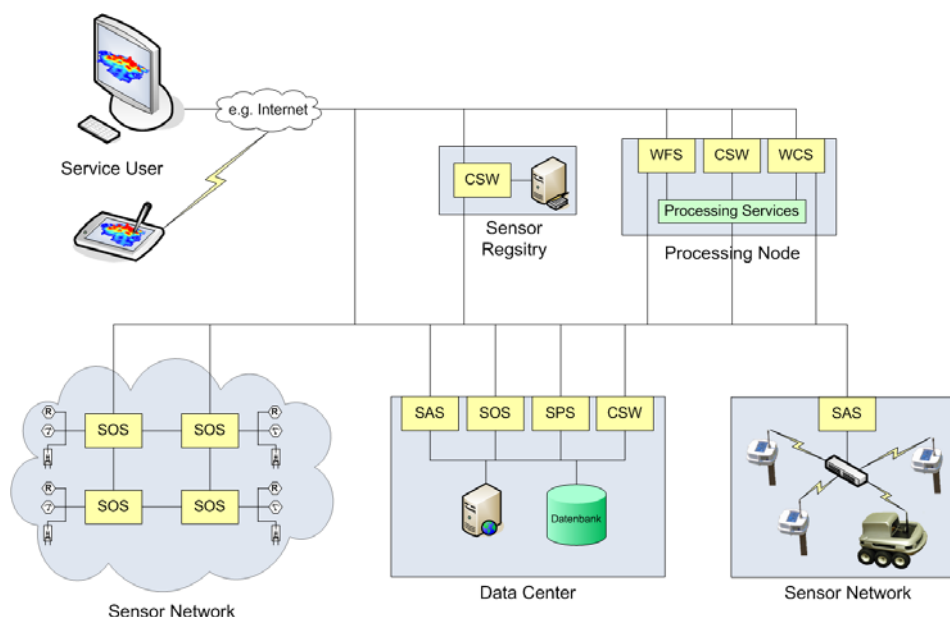


Figure 2 Functional Connections Between the SWE Standards.

It shall be mentioned that the OGC is currently establishing the so-called *SWE Common* namespace specification, which aims at grouping elements, which are used in more than one standard of the SWE family. In effect, this will minimise redundancy, and optimise re-usability and efficiency of the standards. *SWE Common* will mainly comprise very general elements such as counts, quantities, time elements or simple generic data representations.

More information on the Sensor Web Enablement initiative, the incorporated standards and the efforts to embed it into the OGC standard service development can be found on the OGC web site¹.

Apart from the SWE standard family, which is used throughout the sensor network process chain, other OGC standards are employed to build up the overall infrastructure. At first, the Web Processing Service (WPS) is used for integrating measurement data with thematic process models in order to generate contextual information layers. WPS (Schut, 2007) basically allows for the implementation and execution of pre-defined analysis processes with dedicated input and output parameters. It supports synchronous and asynchronous data processing in order to enable sophisticated processing of large amounts of vector and raster data.

For standardised data provision and visualisation, the OGC developed a set of service interfaces dealing with various kinds of GIS data types. The Web Feature Service (WFS), the Web Map Service (WMS) and the Web Coverage Service (WCS)

¹ <http://www.opengeospatial.org>

standards allow for access to data representations such as vectors (point, line, polygon), raster images, and coverages (surface-like structures). More information about these standards and service implementations can be found on the OGC web site¹.

A major advantage of the OGC processing and data provision services mentioned above is the wide variety of standardised (GML, KML etc.) and custom output formats (GeoRSS, PDF etc.). This allows for the integration of the OGC service outputs into other processing, visualisation or decision support services including legacy COTS and open-source GIS analysis software.

3.3 Sensor Network

Within *Common Scents*, the CitySense sensing network (Murty et al., 2008) was used as the underlying sensing and data collection infrastructure. The main goal of the ongoing CitySense project is to build 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. Currently, the network consists of 16 nodes deployed around the city of Cambridge measuring different environmental parameters such as CO₂ concentrations, air temperature, wind speed and direction, or precipitation. The final CitySense deployment will comprise up to 100 sensing nodes.

In order to achieve an interoperable and flexible sensing infrastructure, we implemented several standardised services on top of the CitySense network, in accordance with the *Real-time Geo-awareness* approach (Resch et al., 2009). For data access, we developed a Sensor Observation Service (SOS), which supplies measurement data in the standardised O&M format. It builds the O&M XML structure dynamically according to measured parameters and filter operations. To generate alerts e.g. in case of exceedance of a threshold, we implemented an XMPP (Extensible Messaging and Presence Protocol) based Sensor Alert Service (SAS). It is able to detect patterns and anomalies in the measurement data and generate alerts and trigger appropriate operations such as sending out an email or a text message, or to start a pre-defined GIS analysis process.

3.4 Event-based Alerting

Event recognition and processing happens in two different stages of the workflow. 1.) at sensor level, Complex Event Processing (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. Thus, it can be considered a mechanism for quality control and error prevention. To be able to detect condition changes in measurement values, we enhanced CEP and Event Stream Processing (ESP) techniques by the *location* parameter. Thus, these methods can also serve for the federal organisation of pre-defined geographical domain violations like geo-fences, and for tracing and analysing spatial patterns. 2.) after the data harmonisation process, CEP serves for spatio-temporal pattern recognition, anomaly detection, and alert generation in case of threshold transgression.

Figure 3 shows the components of the CEP-based event processing component, which is built up in a modular structure.

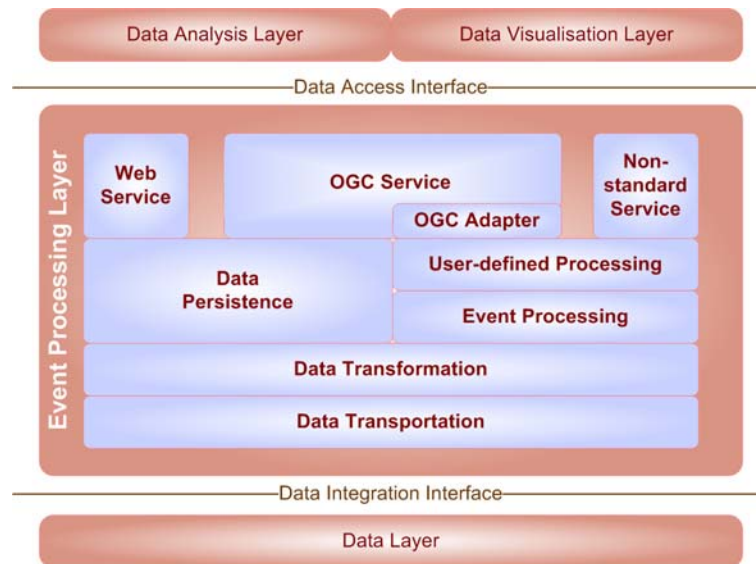


Figure 3 CEP-based Event Processing Component Architecture.

Basically, the event processing component serves as a connection between the data layer (sensor measurements) and the data analysis and data visualisation components, i.e. it prepares raw data in order to be process-able in the analysis and the visualisation layers. The first module is the data transportation, which connects different real-time and non-real-time data sources, i.e. it serves as an entry point into the event processing layer. Next, the retrieved data is passed on to the data transformation module, which prepares the data to be further processed. This “processing” shall not be seen as data analysis, but more as data preparation. Basically, the transportation module converts the byte input stream to objects. These objects can then be used by two higher-level components; firstly, by the data persistence component, which establishes a static data structure from the source data; secondly, by an event processing engine, which processes a real-time stream, identifies/selects events and pushes them to the user-defined processing module. The latter performs a kind of “event filtering” and sends the resulting message to the service components.

One particularity, which shall be mentioned at this point, is the connection between the processing and the data persistence modules. The idea behind this functional link is that data, which have been prepared by the processing components, can either be pushed to the service components or they can be temporarily stored to be accessed by OGC services.

The two service-related components in Figure 3 (Web Service and Non-standard Service) serve as the direct interfaces to the data integration and data analysis layers. They offer the pre-processed raw data via a defined data structure, e.g. in a standardised output format such as GML, KML, geoTIFF etc. or in a custom output format.

For the OGC service component, all data are served over well-established open and standardised interfaces (OGC WFS, WMS and WCS). These XML web interfaces enable standardised data access and guarantee combinability of the various kinds of used data for further automated processing, as described in sub-chapter 3.2. In this way, pre-defined aggregation services can be implemented in the data analysis layer

offering the results to a range of different users, i.e. platforms.

In addition to the standardised interfaces, also a non-standard service has to be created as existing OGC services don't support push mechanisms per se. A longer term option will be to replace these non-standard interfaces by push-capable standard services.

3.5 Real-time Data Harmonisation

During the last years, we faced a drastic increase in available GIS data sources, and this applies especially to rapid developments and price reduction in sensing technologies. To make use of this immense amount of data within environmental monitoring systems, real-time data integration mechanisms have to be developed, which harmonise and fuse the different kinds of data. Furthermore, these data have to be provided in standardised formats in order to allow interoperability and collaboration between different institutions.

Currently, most data integration systems make use of a temporary database to combine different kinds of raw data, as stated in chapter 2. This approach has two distinctive disadvantages. Firstly, it manifests data into a physical structure and thus severely limits real-time capabilities. Secondly, the laborious operation of creating and filling a database table adds another step in the overall workflow, which decreases performance and expands implementation costs and complexity.

To overcome these shortcomings, we implemented the real-time data integration component in a custom datastore for the open-source server GeoServer. This solution offers two main advantages: at first, data are fused on-the-fly in a highly dynamic, fast and parallelised process. At second, GeoServer provides standardised WFS, WMS and WCS outputs, as described above, which allows for simple integration into analysis and visualisation software. More about implementation details can be found in Mittlboeck and Resch (2008).

3.6 Spatio-temporal Analysis

After having set up the infrastructure described above, we implemented two GIS analysis mechanisms. They are both based on ESRI's ArcGIS software suite as reliable and comprehensive open-source GIS processing service implementations do not exist yet.

The first analysis component performs a spatial Inverse Distance Weighting (IDW) interpolation on temperature measurements, which will be used in further research efforts for correlation operations with emission distribution or traffic emergence, and for the detection of urban heat islands. The output of the interpolation process – a navigable 3D map – is shown in Figure 4.




Figure 4 Spatial Temperature IDW Interpolation.

The second processing component uses ArcGIS's 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, we used CO2 data captured by the CitySense network in Cambridge. This 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 5 shows the temporal gradient of the measurement values. Running the time series then 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 a next step.

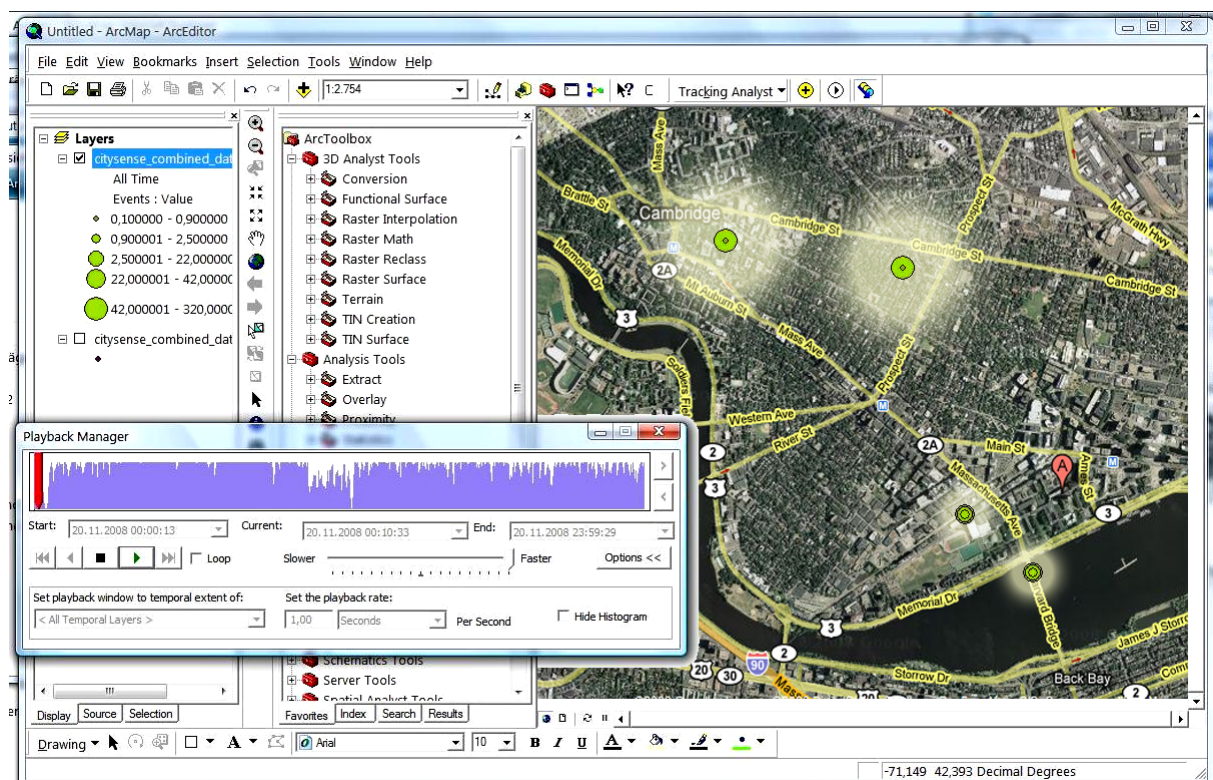


Figure 5 Spatio-temporal CO2 Data Analysis Using Tracking Analyst.

4. POSSIBLE APPLICATION AREAS

Within the *Common Scents* project, we focus on the use case of air quality monitoring for use in the public health sector. However, we designed the monitoring infrastructure in such a modular way that it is not bound to one single application area. Below, several practically motivated fields of real-world applications are described, which could use the same infrastructure presented in chapter 3.

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.

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 (UFPs), particles that are less than 0.1 microns (100 nm) 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.

The use case of *noise mapping* 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. These disputes also arise 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 *Common Scents* infrastructure.

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 Alliance: <http://www.cambridgeenergyalliance.org>) 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.

5. PARTICULAR NON-TECHNICAL CHALLENGES FOR URBAN SENSING

Apart from well known general monitoring challenges, very particular issues arise for 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.

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.) 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 are very general and well elaborated and shall not be discussed in this paper.

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 6 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 influences 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., 2008).




Figure 6 Feedback Loop for Environmental Monitoring in the City.

The first and probably 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 a surrogate 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.

Another challenge is 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.

In consequence, 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 environmental 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, there are some more trivial or obvious challenges to meet, such as optimal positioning of sensor, high spatial and temporal variability of measured parameters or rapid changes in the urban structure, which might cause considerable bias in the measurements.

6. CHANGING THE PERCEPTION OF THE CITY

Generally speaking, fine-grained urban sensing greatly enhances our knowledge of the environment 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. 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 chapter 5.

In terms of public health, the primary impact of the *Common Scents* project 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.

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?

From a quantitative viewpoint, the air quality data measured with the *Common Scents* 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 in the US 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.

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. We claim that this should 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.

7. CONCLUSION

Ubiquitous and continuous environmental monitoring is an enormous challenge, and this is particularly true in the urban context, which poses very specific challenges as well technologically as socially and politically. In this paper we discussed these issues, and outlined how our approach can meet resulting requirements.

Based on the *Real-time Geo-awareness* approach, we developed the *Common Scents* concept, which tries to establish an interoperable, modular and flexible sensing and data analysis infrastructure, as opposed to hitherto monolithic sensor networks. To prove our system's portability, we did implementations in different application domains (public health, environmental monitoring, biometric patient surveillance) using the same underlying infrastructure.

We also elucidated that the impact of pervasive sensing in the city has to be carefully assessed. We 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 market, the insurance sector or urban planning in general.

As the *Common Scents* concept has been developed and implemented together with the Public Health Department of the City of Cambridge, MA US as a concrete end user, we believe that our approach can respond to dedicated needs of the city management. Therefore, the longer-term goal is to enhance people's perception of their environment by adding unseen information layers and thus changing their short-term behaviour by providing real-time decision support.

ACKNOWLEDGEMENTS

The *Common Scents* project is a concerted effort between the Research Studio iSPACE, the MIT SENSEable City Lab, the City of Cambridge's Public Health Department and the Harvard University School of Engineering and Applied Sciences. We would like to thank all internal and external collaborators for making this project happen.

REFERENCES

a) Journal Papers

Balazinska, M., Deshpande, A., Franklin, M.J., Gibbons, P.B., Gray, J., Hansen, M., Liebhold, M., Nath, S., Szalay, A., Tao, V. (2007) Data Management in the Worldwide Sensor Web. **IEEE Pervasive Computing**, vol. 6, no. 2, pp. 30-40, Apr-Jun, 2007.

Kansal, A., Nath, S., Liu, J. and Zhao, F. (2007) SenseWeb: An Infrastructure for Shared Sensing. **IEEE Multimedia**, vol. 14, no. 4, pp. 8-13, October-December 2007.

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**, vol. 19, no. 8-9, pp. 957-973.

Paulsen, H. and Riegger, U. (2006). SensorGIS - Geodaten in Echtzeit“. In: **GIS-Business**, vol. 8/2006, pp. 17-19, Cologne.

Paulsen, H. (2008) PermaSensorGIS – Real-time Permafrost Data. In: **Geoconnexion International Magazine**, vol. 02/2008.

Riva, O. and Borcea, C. (2007) The Urbanet Revolution: Sensor Power to the People!. **IEEE Pervasive Computing**, vol. 6, no. 2, pp. 41-49, Apr-Jun, 2007.

Wood, J., Dykes, J., Slingsby, A. and Clarke, K. (2007) Interactive Visual Exploration of a Large Spatio-temporal Dataset: Reflections on a Geovisualization Mashup. **IEEE Transactions on Visualization and Computer Graphics**, vol. 13, no. 6, pp. 1176-1183, Nov/Dec, 2007.

b) Papers Presented in Conferences

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.

Mittlboeck, M. and Resch, B. (2008) Standardisierte Integration von Real-time Sensormessungen für Zeitnahe GIS-Analyse. In: Strobl, J., Blaschke, T., Griesebner, G. (Eds.): **Angewandte Geoinformatik 2008**, Wichmann Verlag, Heidelberg, pp. 112-117.

Rahm, E., Thor, A. and Aumueller D. (2007) Dynamic Fusion of Web Data. **XSym**

2007, Vienna, Austria, p.14-16.

Resch, B., Calabrese, F., Ratti, C. and Biderman, A. (2008) An Approach Towards a Real-time Data Exchange Platform System Architecture. **Sixth Annual IEEE-IARIA International Conference on Pervasive Computing and Communications**, Hong Kong, 17-21 March 2008.

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. Proceedings of the **IEEE International Conference on Advanced Geographic Information Systems & Web Services - GEOWS2009**, 1-7 February 2009, Cancun, Mexico, pp.92-97.

c) Other Documents

Botts, M., Robin, A., Davidson, J. and Simonis, I. (eds.) (2006) OpenGIS Sensor Web Enablement Architecture. **OpenGIS Discussion Paper**, version 1.0, OGC 06-021r1, 4 March 2006.

Center for Coastal and Land-Margin Research (2009) CORIE. <http://www.ccalmr.ogi.edu/CORIE>. (14 January 2009)

Harrie, L. (2004) Using Simultaneous Graphic Generalisation in a System for Real-Time Maps. **Papers of the ICA Workshop on Generalisation and Multiple Representation**, August 20-21, 2004, Leicester (electronic version available at <http://ica.ign.fr/Leicester/paper/Harrie-v2-ICAWorkshop.pdf>).

Nath, S., Liu, J. and Zhao, F. (2006) Challenges in Building a Portal for Sensors World-Wide. **First Workshop on World-Sensor-Web: Mobile Device Centric Sensory Networks and Applications (WSW'2006)**, Boulder CO, Oct 31, 2006.

National Oceanic and Atmospheric Administration (2008) nowCOAST: GIS Mapping Portal to Real-Time Environmental Observations and NOAA Forecasts. <http://nowcoast.noaa.gov>. (15 December 2008)

Sarjakoski, T. et al. (2004) Geospatial Info-mobility Service by Real-time Data-integration and Generalisation. <http://gimodig.fgi.fi>. (22 December 2008)

Rittman, M. (2008) An Introduction to Real-Time Data Integration. <http://www.oracle.com/technology/pub/articles/rittman-odi.html>, 14 June 2008.

Schut, P. (ed.) (2007) Web Processing Service. **OpenGIS Standard**, version 1.0.0, OGC 05-007r7, 8 June 2007.

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.

Sybase Inc. (2008) Real-Time Events Data Integration Software. <http://www.sybase.com/products/dataintegration/realtimeevents>. (16 December 2008)

University of Oklahoma (2009) OKCnet. <http://okc.mesonet.org>. (12 February 2009)