

Report from the First Workshop on Geo Sensor Networks

S. Nittel,* A. Stefanidis,† I. Cruz, M. Egenhofer, D. Goldin, A. Howard, A. Labrinidis
S. Madden, A. Voisard, M. Worboys
Portland, ME, Oct 9-11 2003
University of Maine, USA

1 Introduction

Advances in sensor technology and deployment strategies are revolutionizing the way that geospatial information is collected and analyzed. For example, cameras and GPS sensors on-board static or mobile platforms have the ability to provide continuous streams of geospatially-rich information. Furthermore, with the advent of nanotechnology it becomes feasible and economically viable to develop and deploy low-cost, low-power devices that are general-purpose computing platforms with multi-purpose on-board sensing and wireless communications capabilities. Today, research efforts are taking place developing infrastructure for systems consisting of large numbers of small unattended, untethered and collaborative sensor nodes that have non-renewable power supply and communicate via short range radio frequency with neighboring nodes. These types of sensors may also act collaboratively within broader network configurations which can range in scale from a few cameras monitoring traffic to thousands of nodes monitoring an ecosystem. The challenge of sensor networks is to aggregate sensor nodes into computational infrastructures that are able to produce globally meaningful information from raw local data obtained by individual sensor nodes.

In geo sensor networks the geospatial content of the information collected, aggregated, analyzed, and monitored by a sensor network is fundamental; analysis and aggregation might be performed locally in real-time by the sensor nodes or between sensor nodes, or off-line in several distributed, in-situ or centralized repositories. Thus, a geosensor network can loosely be defined as a sensor network that monitors phenomena in geographic space. Geographic space can range in scale from the confined environment of a room to the highly complex dynamics of an ecosystem region.

The spatial aspects of the overall technology is important on multiple (abstraction) levels of a geo sensor net-

work, as the concepts of space, location, topology, and spatiotemporal events are modelled on various abstraction levels. For example, the hardware and communication layers handle the physical space of sensor deployment, and communication topologies. The database layer generates execution plans for spatiotemporal queries that relate to sensor node locations and groups of sensors. Applications deal with the relation between sensor networks and phenomena in geographic space. We feel that the academic and practical expertise of the spatial information theory and engineering domain are crucial to advance the development of sensor networks on all different abstraction levels. The ultimate objective is to develop generic sensor network programming infrastructure that is reusable, and widely applicable in the different application domains types.

2 Workshop

The first Geo Sensor Networks workshop took place in Portland, Maine, Oct 9-11 2003, and was co-organized by Silvia Nittel and Anthony Stefanidis, both from the National Center of Geographic Information and Analysis, University of Maine. Thirty-two researchers from diverse research domains attended the workshop, presenting papers, and participating in panel discussions¹.

2.1 Purpose

The purpose of this workshop was to bring together researchers from the areas of (spatial) database management systems and spatial information modelling, as well as operating systems, robotics, mobile computing, image analysis, and environmental applications to provide a discussion forum for experts who are interested in developing infrastructure for and being users of sensor networks. We expected that the different expertise regarding spatial information modelling and handling found in the different areas

*nittel@spatial.maine.edu, NCGIA, University of Maine

†tony@spatial.maine.edu, NCGIA, University of Maine

¹The GSN URL is <http://www.spatial.maine.edu/gsn03/>

would provide a thought-provoking theme for all participants to successfully deal with the challenges of developing and deploying geo sensor networks.

In detail, the workshop addressed issues related to the collection, management, analysis, processing and delivery of real-time geospatial data streams, mobile computing and context-aware computing, temporal-spatial queries over geo sensor networks, sensor data mining, higher level abstraction for modelling sensor network data, and virtual reality modelling.

2.2 Structure

The workshop was structured into a combination of research paper presentations, invited keynotes and panel sessions. The research papers were submitted following a call for papers, and were peer reviewed. It was also considered valuable to stimulate workshop interaction through invited keynotes by some of the leading experts in the fields of sensor networks and geospatial data modelling (Samuel Madden - UC Berkeley, Agnes Voisard - Fraunhofer ISST and FU Berlin, Alexandros Labrinidis - U Pittsburgh, Max Egenhofer - UMaine, and Mike Worboys - UMaine). A book including the original research contributions and invited papers will be published by the CRC division of Taylor and Francis in early 2004.

3 Workshop Highlights

This section includes a summary of the workshop highlights that cut across several keynotes and paper presentations.

3.1 Programming Sensor Networks using DBMS Technology

The first highlight of the workshop was centered around assessing the state of the art in sensor network prototype implementation as presented by the invited keynotes of Samuel Madden (UC Berkeley) [1] and Alex Labrinidis (UPittsburgh) [2]. In general, in the database community the assumption is made that programming sensor networks is hard, and database management system (DBMS) technology with its characteristics of declarative data models, query languages and automatic query optimization makes the job of programming sensor networks significantly simpler. DBMS-style query execution over sensor networks is developed with the requirement that queries are formalized in such a way that their execution plans over the sensor network infrastructure are automatically optimizable by the DBMS. Hereby, the main optimization criteria is energy-efficient processing of information since batteries are typi-

cally not renewed during the lifetime of an application deployment. Since the transmission of data between sensor nodes is costly with regard to energy consumption, optimization attempts to minimize communication between nodes while guaranteeing quality of service. Strategies include minimization of data acquisition, i.e. instructing sensor nodes to only generate (sample) the data that is necessary for a query, or to only forward new values that are within a significant threshold change of the current sampling values. Another strategy is to exploit automatic operator reordering during query processing so that operators that are 'cheaper' (i.e. lower drain on energy to obtain a sensor sample) are evaluated first, and sampling of more 'expensive' sensors for a conjunctive predicate can be avoided. Other strategies are compressing values so that less data is transmitted between nodes, or suppressing values within a temporal coherency tolerance.

Today, power consumption is driven by sampling sensor values, and listening to queries. Minimizing the listing time of sensor nodes allows them to only wake up and synchronize for very short periods of time. With such a massively distributed computing system the notion of synchronized system time is a major challenge. Also, sampling frequency can be adapted over time to prolong the battery life time of sensor nodes.

3.2 Scale and Mobility of Sensor Nodes

Scale of sensor data collection and processing was identified as a significant challenge in geo sensor networks. Varying scales of sensor data collection and processing are required for different aspects of a problem or even a particular user. The issue matters with regard to sensor node locations and their distribution density, the size of regions of interest, and intervals of sampling. Also, user and application needs play a significant role as such to collect raw data, statistical data, or models, and the level of quality of service such as freshness of data, response time, etc.

To enable multi-resolution queries, different epoch sizes can be assigned to different spatial areas of the network. Shorter epochs enable a higher frequency data sampling and aggregation. Another alternative consists of a group-based routing tree construction. A 'group' is a set of sensors that e.g. exhibits the same capabilities (e.g. temperature sensing), and the routing tree consists of parent-child nodes of the same group while all nodes are collocated. This decreases the number of messages a parent node has to send, and the number of queries to respond to. Simulation results demonstrate that this mechanism works well for a small number of different groups, but a larger number of members per group.

For today's prototypes, the assumption is made that sen-

sensor nodes are stationary for the time being. However, it is most likely that sensors are mobile by either being self propelled or being attached to moving objects. In the environmental domain, sensors might be floating in a drainage or be carried by the wind in storms. Network protocols contain built-in mechanisms to construct flexible routing trees despite the mobility of sensor nodes. Nevertheless, sensor nodes need to be able to geolocate their own position with sufficient accuracy, a problem that is still open today. Current research work in robotics with regard to self localization of robots could be leveraged [6]. Likely, sensors nodes are rarely located at exactly the position that is necessary for a spatial region query in the geographic space. Mappings between higher-level spatial user predicates and actual physical sensor node locations are of interest, and also constructing an optimal routing tree for a specific spatial query predicate [3]. Furthermore, the density of sensor nodes needs to be mapped to different application resolution needs. Dense deployment of sensor nodes is economically not viable. Mechanisms such as robots fixing density problems by 'dropping' sensor nodes in low density areas might be a more flexible and economic solution.

3.3 Higher-level Modelling and Reasoning

So far, higher-level data models and query languages that allow reasoning over the data collected via sensor networks and express complex interactions are not available. This will be necessary to more fully exploit the information that the network can provide to explain and make predictions about the domains in which the sensors are embedded. Invited talks from Max Egenhofer and Mike Worboys (University of Maine) discussed some of the issues around the provision of higher-level modelling and reasoning capabilities.

A field model was proposed in which distributions of spatial attributes, along with their sampling and interpolation protocols could be formally described [4]. A field model allows to formally define objects such as 'toxic cloud' or a temperature field over a certain area whereby the underlying sample points are created by sensor nodes. Values at other than sensor node locations are interpolated via operators of the field model. The dynamic nature of the world also leads to richer modelling and querying abilities than are usual for spatial databases. A proposal was made to relate the work on formal models of computational processes to real world event models. Sensor networks monitor collections of occurrences as mentioned above. Occurrences can be either processes ('the car is moving fast'), or events ('the car stops at the lights'). Occurrences also relate to physical objects; in this example the car is the physical object. It is necessary to capture and relate occur-

rences and physical objects in a more formal way. An ontology of events, processes and actions will form the foundation of sensor-based models of the dynamic world. In order to fully exploit such an event-driven approach, event properties, event-event relationships, and the ways that objects and fields can participate in events needs analysis [5]. Other important related issues include the abstraction and summarization of meaningful entities in dynamic domains, the role of triggers, event notification systems, along with the ability to map these higher-level concepts onto the sensor database model and query architecture [7]. For humans to fully exploit the power of sensor networks, cognitive issues related to distributed computational processes and event-based models need to be explored. In particular, work is required on human interaction with sensor networks.

3.4 The Geo Sensor Web

Today, many sensor field station are already in place in the environmental domain. Most of the information is collected locally at the stations, and often retrieved manually. With the advent of sensor network hardware, more of the sensor data will be available online and in realtime, and sensor networks will likely be integrated with existing field stations. The question is whether the computational paradigms of ad-hoc collaborations between sensor nodes will extend to heterogeneous types of sensors, and thus, an ad-hoc 'sensor web' similar to the world wide web is the future paradigm of sensor networks. We can expect that real-time sensor data is accessible from everywhere at any time, and can be combined in new ways with little programming effort. To enable such a 'pervasive sensor network infrastructure', interoperability protocols for sensor networks are a key issue. Other aspects are discovery of sensor data sources, and meta data for sensor data streams, etc.

3.5 Sensor Networks Enabling Virtual Geo Reality

The development of realistic virtual reality (VR) models of urban environments has been the topic of substantial research efforts in the last few years (see e.g. the Virtual LA project at UCLA). These VR models of urban scenes are photorealistic: they provide views of the world very similar to the ones we would perceive if we were to roam the scene, sometimes even to the point of including graffiti on the walls. However, these models are not tempo-realistic: the real world is in flux, yet these models represent only a single instance of the scene, namely the moment when the images used to create them were actually collected. Consid-

ering the high cost to actually build such models, their updating is rarely a priority, unless of course specific information (e.g. the demolition of an important building) makes it necessary to update a small part of the database. Furthermore, it is often remarked that VR models feel empty, failing to incorporate the movement of vehicles and people. This lack of temporal validity has hindered the use of virtual models as convenient interface to spatial databases, even though they convey geospatial information and their expressive power is of tremendous value to the communities that use geospatial information in everyday activities.

Geosensor networks enable the evolution from VR to Virtual GeoReality (VGR) models, offering spatiotemporally accurate models of reality. Now, VGR models can be suitable to monitor and communicate current and emerging situations by enhancing them with on the fly update capabilities, and the ability to monitor and model evolving activities. On the other hand, a need exists to visualize realtime sensor network data to deal with cognitive aspects of humans learning to use, program, and monitor sensor networks. Towards this goal, change detection and monitoring dynamic phenomena are becoming important research issues in image processing. Workshop presentations addressed certain issues related to this transition towards VGR models, mostly related to object tracking using distributed video sensors.

4 Outlook and Open Issues

Geo sensor networks are a rapidly evolving multidisciplinary field that challenges the research areas involved to integrate new techniques, models and methods that are often not found in their classical research agendas. Interdisciplinary workshops like the first Geo Sensor Networks meeting are an important step towards providing an exchange forum for this newly emerging community. Due to the large overlap of research challenges but varying backgrounds in the different domains, such workshops can be a fruitful opportunity for collaborations. During several panel discussions, open issues were discussed.

One of the prominent open issues using sensor networks today is the issue of sensor data *privacy*. With the requirements to design ultra-light wireless communication protocols for small-form devices not much room is left for advanced encryption schemes. A related issue is the need for *authentication* of sensed data. If sensor networks are deployed in security sensitive areas, built-in mechanisms need to be available to provide for such data authentication. A third open issue is *data quality*. Mechanisms need to assure that defective or incorrectly calibrated sensors are excluded from the computation, and that calibration is established individually as well as collectively before deploy-

ment and also continuously later on. Today, many research efforts in sensor networks are conducted under assumptions derived from the constraints of current hardware platforms such as the Berkeley motes. Many of these assumptions such as using radio broadcasting as communication modality or restricted battery life might not be valid anymore in a few years, and these assumptions might change completely.

A follow-up workshop is planned for Fall 2004.

Acknowledgements

This workshop was supported by the National Science Foundation under NSF grant number EPS-9983432.

References

- [1] Madden, S. *Query Processing Challenges in Sensor Networks*, Invited Talk, First Workshop on Geo Sensor Networks, Portland, Maine, October 9-11 2003.
- [2] A. Labrinidis, and P. Chrysanthis, *Location-Aware In-Network Data Aggregation for Sensor Networks*, Invited Talk, First Workshop on Geo Sensor Networks, Portland, Maine, October 9-11 2003.
- [3] D. Goldin, H. Dave, A. Kutlu, M. Song, *Geo-routing and Delta-gathering: Efficient Data Propagation Techniques for GeoSensor Networks*, First Workshop on Geo Sensor Networks, Portland, Maine, October 9-11 2003.
- [4] M. Egenhofer, *Sensors, Sensors Everywhere*, Invited Talk, First Workshop on Geo Sensor Networks, Portland, Maine, October 9-11 2003.
- [5] M. Worboys, *Event-based Models of Geosensor Networks*, Invited Talk, First Workshop on Geo Sensor Networks, Portland, Maine, October 9-11 2003.
- [6] A. Howard, G. Sukhatme, M. Mataric, *From Mobile Robot Teams To Sensor and Actuator Networks*, Invited Talk, First Workshop on Geo Sensor Networks, Portland, Maine, October 9-11 2003.
- [7] A. Voisard, *Information Delivery and Exchange in Mobile Applications - Concepts and Applications in Tourism*, Invited Talk, First Workshop on Geo Sensor Networks, Portland, Maine, October 9-11 2003.