

05381 Abstracts Collection
Form and Content in Sensor Networks
— Dagstuhl Seminar —

Leonidas J. Guibas¹, Uwe D. Hanebeck² and Thomas C. Henderson³

¹ Stanford Univ., US
guibas@cs.stanford.edu

² Univ. Karlsruhe, DE
hanebeck@ira.uka.de

³ Univ. of Utah, US
tch@cs.utah.edu

Abstract. From 18.09.05 to 23.09.05, the Dagstuhl Seminar 05381 “Form and Content in Sensor Networks” was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Sensors, signal processing, sensor networks, intelligent systems, sensor data processing

05381 Executive Summary – Form and Content in Sensor Networks

From the September 18th until September 23rd, 2005 a Dagstuhl Seminar took place with the topic "Form and Content in Sensor Networks". 26 participants from four different countries, which are experts in sensor networks from the topics information processing, communication and robotics, presented current state of the art in the field of algorithm for sensor networks and how content and structure impact information processing in the networks. The presentations ranged from very theoretical computational models and algorithms to prototype implementations for monitoring the environment.

Keywords: Sensors, signal processing, sensor networks, intelligent systems, sensor data processing

Joint work of: Guibas, Leonidas J.; Henderson, Thomas C.; Hanebeck, Uwe D.

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2006/754>

Design of an Artificial Swarm: The I-SWARM Project

Natalie Bender (Universität Karlsruhe, D)

In classical microrobotics, highly integrated and specialised robots have been developed in the past years, which are able to perform micromanipulations controlled by a central high-level control system. On the other hand, technology is still far away from the first "artificial ant" which would integrate all capabilities of these simple, yet highly efficient swarm building insects.

This has been the motivation of other research fields focusing on studying such swarm behaviour and transferring it to simulation or physical robot agents. Realisations of small robot groups of 10 to 20 robots are capable to mimic some aspects of social insects, however, the employed robots are usually huge compared to their natural counterparts, and very limited in terms of perception, manipulation and cooperation capabilities. The project aims to take a leap forward in robotics research by combining experts in microrobotics, in distributed and adaptive systems as well as in self-organising biological swarm systems. The project aims at technological advances to facilitate the mass-production of microrobots, which can then be employed as a "real" swarm consisting of up to 1,000 robot clients. These clients will all be equipped with limited, pre-rational on-board intelligence. The swarm will consist of a huge number of heterogeneous robots, differing in the type of sensors, manipulators and computational power. Such a robot swarm is expected to perform a variety of applications, including micro assembly, biological, medical or cleaning tasks.

Building on a large expertise in microrobot technologies, the project addresses topics like polymer actuators, collective perception, using (instead of fighting) micro scaling effects, artificial and collective intelligence. The project results will enable humans to further understand the micro world, bridge the gap between micro and nano technologies and be the stepping stone to a "real artificial ant".

Keywords: Swarm robotic, Self organization

Joint work of: Bender, Natalie; Wörn, Heinz; Seyfried, Jörg

Full Paper:

<http://microrobotics.ira.uka.de/>

Sensor Fusion for Diurnal Scene Understanding

Bir Bhanu (Univ. California - Riverside, USA)

This talk will discuss our current work on dynamic sensor fusion using physical and statistical models of sensors that are used in an evolutionary computational framework for dynamic scene understanding in a network of sensors.

Organic Middleware for Self-Organizing Sensor Networks

Uwe Brinkschulte (Universität Karlsruhe, D)

The Organic Computing paradigm proposes to build complex systems more life-like. Such systems should have features found in organic entities like self-organization, self-configuration, self-optimization, self-healing, self-protection, self-awareness, etc. Since sensor networks are complex systems, it sounds natural to apply this paradigm here. Many of the organic features mentioned above can be supported by middleware, which interconnects the sensor nodes and organizes task distribution. So we propose an architecture for an organic middleware for sensor networks. This architecture contains:

- *An hierarchical self-organisation of the sensor nodes. Therefore, sensor nodes are self-aware of their own capabilities and order themselves in hierarchical levels and groups. To scale to large number of nodes, group leaders are elected according to node capabilities.

- *Self-optimization is performed by constantly monitoring and evaluating critical system parameters like energy consumption or run-time efficiency. A task redistribution and a re-election of leaders is done according to evaluation results.

- *Self-healing is done in a similar way to self-optimization. In case of node or communication failures, tasks are redistributed and leaders are re-elected if necessary.

- *To allow a unique middleware structure for all hierarchy levels, a nanokernel approach is proposed. This nanokernel is suitable for simple sensor nodes with little resources and can perform only a very basic functionality and restricted services. For more complex nodes, the nanokernel is expanded to a microkernel by adding a specific service. The microkernel allows improved functionality and flexible service handling. For full size nodes, the microkernel can be supplemented now by further specific services to perform requested tasks.

- *Special algorithms have to be developed to establish (at least) soft real-time capabilities for the sensor network in the context of self-organisation.

Keywords: Self-Organization, Organic Computing, Middleware

Joint work of: Brinkschulte, Uwe; Pacher, Mathias; Picioroaga, Florentin

Algorithmic aspects of large sensor networks

Sándor Fekete (TU Braunschweig, D)

We present a new framework for the crucial challenge of self-organization of a large sensor network. The basic scenario can be described as follows: Given a large swarm of immobile sensor nodes that have been scattered in a polygonal region, such as a street network. Nodes have no knowledge of size or shape of the environment or the position of other nodes. Moreover, they have no way of

measuring coordinates, geometric distances to other nodes, or their direction. Their only way of interacting with other nodes is to send or to receive messages from any node that is within communication range. The objective is to develop algorithms and protocols that allow self-organization of the swarm into large-scale structures that reflect the structure of the street network, setting the stage for global routing, tracking and guiding algorithms.

Our algorithms work in two stages: boundary recognition and topology extraction. All steps are strictly deterministic, yield fast distributed algorithms, and make no assumption on the distribution of nodes in the environment, other than sufficient density.

Keywords: Distributed algorithms, sensor networks, boundary recognition, topology extraction

Joint work of: Fekete, Sándor; Kröller, Alexander; Pfisterer, Dennis; Fischer, Stefan

Full Paper:

<http://arxiv.org/abs/cs.DC/0510048>

See also: To appear in: SODA 2006, paper available at arxiv.org

Distributed Perception Networks, a MAS approach to robust and efficient fusion of heterogeneous information.

Frans Groen (University of Amsterdam, NL)

Distributed Perception Networks (DPN) are Multi Agent Systems (MAS) for robust and efficient situation assessment, where the relevant information sources are not known prior to the operation and the reasoning about a few hypotheses is based on large quantities of very heterogeneous and noisy information. Such situation assessment is indispensable for real world decision making and controlling processes. For example, crisis management focusing on chemical disasters requires quick discovery and localization of toxic gases. Similarly, Tsunami warning systems must be able to identify a potentially dangerous situation by evaluating spatially dispersed information sources, such as buoys, seismic data, human reports, etc. DPNs implement a logical layer on top of existing sensory, communication and processing/storage infrastructure. DPN agents can integrate heterogeneous, spatially dispersed information sources (e.g. sensory systems, databases, GSM networks and www) into meaningful inference systems. A DPN is essentially an organization of agents which cooperate in order to discover the relevant information sources at runtime and establish information filtering channels that support efficient mapping between large quantities of observations and hypotheses about events that cannot be observed directly. At the lowest level different agents with a direct access to information sources interpret raw data while at the higher levels agents use local world models for high-level information fusion. Such local world models are encoded through Bayesian Networks

(BN) and represent basic modeling building blocks. At runtime DPN agents organize autonomously and integrate local Bayesian Networks into complex causal models, which provide mapping between different symptoms and hypotheses. In other words, DPN agents can assemble complex hierarchical classifiers at runtime. Bayesian Networks facilitate fusion of noisy information as well as the distribution and sharing of fusion processes. By distributing the fusion throughout organizations of agents, we can efficiently process large quantities of information and reduce the danger of a single point of failure. In addition, through the DPN modularity the design and maintenance of fusion systems is simplified. Namely, complex models can be assembled out of smaller and simpler partial world models which can be obtained from different experts or machine learning processes.

Moreover, DPNs support accurate reasoning based on very noisy information and subjective modeling parameters. This is very relevant for the real world applications, where we often cannot obtain precise models and information sources are not perfect. We show that DPNs can form distributed BNs that are inherently robust w.r.t. the modeling parameters and facilitate localization of modeling parameters that are inadequate in a given situation. Thus, we can estimate the fusion quality and signal potentially misleading results.

Keywords: Distributed Perception Networks, robustness, self-organization, fault localization

Joint work of: Groen, Frans; Pavlin, Gregor; Maris, Marinus; Nunnink, Jan

Routing and Information Discovery in Sensor Networks

Leonidas Guibas (Stanford University, USA)

As we transition from sensor networks that are used as passive data collection devices to more proactive networks that perform in-network processing and reasoning to task specific sensors, new problems and issues arise. Such systems require point-to-point connections between sensor nodes, so appropriate node naming and routing mechanisms must be developed. We discuss naming and routing schemes that do not require node localization. Nodes name themselves based on their position relative to other nodes in the network. Two such schemes are presented, a two-level method (Glider) based on landmarks, and one based on a hierarchical decomposition of the network. The latter has the property that the routes produced are at most four times as long as the optimal ones; this is a worst-case ratio and in practice the approximation ratio is very close to 1.

We also discuss information brokerage mechanisms that allow node that want information (consumers) to locate node that have the desired information (producers). One such scheme uses Glider and the idea of double rulings. The other uses the hierarchical decomposition and geographic hashing. Both exhibit good load balancing properties.

Keywords: Naming, routing, information brokerage

Verification and Validation of Sensor Networks

Thomas C. Henderson (University of Utah, USA)

Sensor networks play an increasingly important role in critical systems infrastructure and should be correct, reliable and robust. In order to achieve these performance goals, it is necessary to verify the correctness of system software and to validate the more broadly defined world and system models. This includes:

- *Physical Phenomena (PDE models, statistical models, etc.),
- *Signals (Equations of state, physical properties, etc.),
- *Sensors (Physics models, noise models, etc.),
- *Hardware (Failure models, power consumption models, etc.),
- *RF (Antenna models, bandwidth, delay, propagation, etc.),
- *Embedded Code (Correctness, complexity, context),
- *Distributed Algorithms (Correctness, concurrency models, etc.),
- *Overall Sensor Network and Environment Models (Percolation theory, wave theory, information theory, simulation, etc.).

We outline some of the V & V issues involved in the various aspects of sensor networks as well as possible approaches to their development and application both in simulation and in operational deployed systems.

Keywords: Models, verification, validation, sensor networks

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2006/753>

Geometric Summaries for Sensor Nets

John E. Hershberger (Mentor Graphics - Wilsonville, USA)

In many applications of sensor nets, results are communicated to the user via a low-bandwidth channel. Raw data must be summarized before communication to the user, while preserving the data features the user cares about. For example, the set of temperature sensors in a sensor field detecting values above some threshold defines a “hot” region. A user is interested in knowing the geometry of the hot region, but communicating the locations of all the hot sensors out of the sensor net would be wasteful. The sensor net itself must perform some data reduction before communicating the results to the user.

This raises the problem of geometric simplification: given a description of some geometric region, compute a summary that approximates the input well but uses much less storage than the original data. The summary should be simple to compute, since its computation will be resource-limited (either in space, time, or both). This problem arises in other contexts as well, such as massive spatial databases and electronic design automation, to name just two.

We consider several schemes for simplification of geometric data. One technique, called “cluster hulls,” approximates the boundary of a collection of sample

points as a union of convex regions. Each convex region is represented using a provably optimal approximation; the overall representation gives good results in practice. A second simplification technique applies to rectilinear polygons, and gives provable-quality approximations using a fixed-size representation. A third technique, still under development, uses the concept of minimum link paths to represent geometric regions whose boundaries are known only approximately.

Joint work of: Suri, Subhash; Shrivastava, Nisheeth; Buragohain, Chiranjeeb; Gandhi, Sorabh

Probabilistic Methods in Robotic Sensor Networks

Alexei Makarenko (University of Sydney, AU)

- * Multi-robot systems as sensor networks
 - * decentralized estimation algorithms for environment monitoring
 - * experimental results with air vehicles
 - * current work with air/ground/human platforms
 - * parallels between decentralized estimation algorithms and graphical models
 - * future directions with graphical models

Keywords: Multi-robot systems, decentralized data fusion, graphical models

Generic Model and Architecture for Cooperating Objects in Sensor Network Environments

Pedro José Marrón (Universität Stuttgart, D)

The complexity and heterogeneity of cooperating object applications in ubiquitous environments or of applications in the sensor network domain require the use of generic models and architectures. These architectures should provide support for the following three key issues: flexible installation, management and reconfiguration of components in the system; optimization strategies whose implementation usually involves the proper management of cross-layer information; and proper adaptation techniques that allow for the self-configuration of nodes and components in the system with minimal human intervention. In this paper, we present one possible instance of such a generic model and architecture and show its applicability using Sustainable Bridges, a sensor network application that requires the analysis of complex sensor data to achieve its goal of effectively monitoring bridges for the detection of structural defects.

Keywords: Sensor networks, model, deployment, middleware architecture

Inference Meta Models: A New Perspective On Belief Propagation With Bayesian Networks

Gregor Pavlin (University of Amsterdam, NL)

Modern situation assessment and controlling applications often require efficient fusion of large amounts of heterogeneous and uncertain information. In addition, fusion results are often mission critical.

It turns out that Bayesian networks (BN) are suitable for a significant class of such applications, since they facilitate modeling of very heterogeneous types of uncertain information and support efficient belief propagation techniques. BNs are based on a rigorous theory which facilitates (i) analysis of the robustness of fusion systems and (ii) monitoring of the fusion quality. We assume domains where situations can be described through sets of discrete random variables. A situation corresponds to a set of hidden and observed states that the nature ‘sampled’ from some true distribution over the combinations of possible states. Thus, in a particular situation certain states materialized while others did not, which corresponds to a point-mass distribution over the possible states. Consequently, the state estimation can be reduced to a classification of the possible combinations of relevant states. We assume that there exist mappings between hidden states of interest and optimal decisions/actions. In this context, we consider classification of the states accurate if it is equivalent to the truth in the sense that knowing the truth would not change the action based on the classification. Clearly, BNs provide a mapping between the observed symptoms and hypotheses about hidden events. Consequently, BNs have a critical impact on the fusion accuracy. We emphasize a fundamental difference between the model accuracy and fusion (i.e. classification) accuracy. A BN is a generalization over many possible situations that captures probability distributions over the possible events in the observed domain. However, even a perfect generalization does not necessarily support accurate classification in a particular situation. We address this problem with the help of the Inference Meta Model (IMM) which describes information fusion in BNs from a coarse, runtime perspective.

IMM is based on a few realistic assumptions and exposes properties of BNs that are relevant for the construction of inherently robust fusion systems. With the help of IMM we show that in BNs featuring many conditionally independent network fragments inference can be very insensitive to the modeling parameter values. This implies that fusion can be robust, which is especially relevant in many real world applications where we cannot obtain precise models due to the lack of sufficient training data or expertise. In addition, IMM introduces a reinforcement propagation algorithm that can be used as an alternative to the common approaches to inference in BNs. We can show that the classification accuracy of this propagation algorithm is asymptotically approaching 1 as the number of conditionally independent network fragments increases. Because of these properties, the propagation algorithm can be used as a basis for effective detection of misleading fusion results as well as discovery of inadequate modeling components and erroneous information sources.

Keywords: Robust Information Fusion, Bayesian Networks, Heterogeneous Information, Modeling Uncertainties

Joint work of: Pavlin, Gregor; Nunnink, Jan; Groen, Frans

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/756>

A Training Method for Source Isolation of Simple Product-Form Plumes by Sensor Networks

Nageswara Rao (Oak Ridge National Lab., USA)

We consider a class of simple, idealized plumes characterized by a product of injection and distance decay terms. The plume propagates with a constant velocity, and its distance term decays exponentially across a planar region. Using a network of three ideal intensity sensors, the difference triangulation method can identify the origin of plume within a specified precision. We consider that the sensors are subject to random, correlated errors with unknown distributions in measuring the intensities corresponding to a single plume. We present a hybrid training method that utilizes the plume equation together with controlled sensor measurements to identify the plume's origin with probabilistic performance guarantees. The training consists of computing, using the measurements, a suitable precision value for the difference triangulation method to account for sensor distributions. We derive a distribution-free relationship between the training sample size and the precision and probability with which plume's origin is identified.

Keywords: Distributed sensor network, plume detection and identification, difference triangulation, training, sample size

On the Structure of Data Sets Observed by Physically Embedded Networks

Sergio Servetto (Cornell University, USA)

In this position paper, we argue that the mainstream view in our community of what makes for problems in sensor networks that are worthy of attention is very lopsided, focusing much more on pure networking and other classical computer systems issues, than on the equally important but often neglected tasks of modeling of distributed signal fields, and on their acquisition, processing, synthesis and control. We show concrete examples of how these are essential aspects that need to be dealt with in any good system design, and we survey experimental work along these lines currently under way in our lab.

See also: Position paper submitted to IEEE Workshop on Embedded Networked Sensors (EmNetS), 2006.

Detecting Cuts in Sensor Networks

Subhash Suri (Univ. California - St. Barbara, USA)

If sensor networks are to act as our remote “eyes and ears,” then we need to ensure that any significant failure (natural or adversarial) suffered by the network is promptly and efficiently detected. Tracking the operational health of the infrastructure is important in any communication network, but it is especially important in sensor networks due to their unique characteristics, and the need to perform this duty with very little overhead.

In this talk we will consider a concrete problem of detecting linear cuts that isolate at least ϵ fraction of the nodes from the base station. We show that the base station can detect whenever an ϵ -cut occurs by monitoring the status of just $O(1/\epsilon)$ nodes in the network. Our scheme is deterministic and it is free of false positives: no reported cut has size smaller than $\epsilon n/2$. Besides this combinatorial result, we also propose efficient algorithms for finding the $O(1/\epsilon)$ nodes that should act as *sentinels*, and report on our simulation results, comparing the sentinel algorithm with two natural schemes based on sampling.

Joint work of: Shrivastava, Nisheeth; Toth, Csaba

Full Paper:

<http://www.cs.ucsb.edu/~suri/pubs.html>

Design of a Query Language for Stochastic Model-Based Sensor Networks

Christian von der Weth (Universität Karlsruhe, D)

To take advantage of database query languages for gaining access to the information of sensor networks turned out to be a very promising approach. A lot of work has been done to enhance common database query languages for the application in traditional sensor networks.

In our work we address the whole new class of Stochastic Model-Based sensor networks. The main concept of this class of sensor networks is the application of an underlying stochastic model for the measured physical phenomena. This model enables the acquisition of data apart from the fixed position of a single sensor and can help to reduce the number of sensors as well as the number necessary sensor readings for the response of a query. Due to the stochastic background used for the physical model, a new sort of queries emerges.

In this work we outline the characteristics of Stochastic Model-Based Sensor Networks and discuss their effects on the design of an appropriate query language. In particular, we focus on the users possibilities in dealing with the new kind of uncertainties and probabilities caused by the stochastic approach.

We also propose some extensions to the most popular database query language SQL in order to cope with both Stochastic Model-Based and traditional sensor networks.

Keywords: Stochastic Model-Based Sensor Networks database query languages

Joint work of: Brunn, Dietrich; Ringwald Fabian; von der Weth, Christian