Report from Dagstuhl Seminar 13492

Geosensor Networks: Bridging Algorithms and Applications

Edited by

Matt Duckham[1], Stefan Dulman[2], Jörg-Rüdiger Sack[3], and Monika Sester[4]

1 **The University of Melbourne, AU,** mduckham@unimelb.edu.au
2 **TU Delft, NL,** s.o.dulman@tudelft.nl
3 **Carleton University, Ottawa, CA,** sack@scs.carleton.ca
4 **Leibniz Universität Hannover, DE,** monika.sester@ikg.uni-hannover.de

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 13492 "Geosensor Networks: Bridging Algorithms and Applications." New geosensor networks technologies have the potential to revolutionize the way we monitor and interact with the world around us. The objective of the seminar was to move closer to realizing this potential, by better connecting theoretical advances with practical applications and education. The Seminar ran from 1–6 December 2013, and brought together 21 participants from around the world, representing wide variety of disciplinary backgrounds and expertise connected with geosensor networks. While these discussions are continuing to develop and bear fruit, this report summarizes the results of the discussions held at the seminar.

Seminar 1.–6. December, 2013 – www.dagstuhl.de/13492
1998 ACM Subject Classification C.2.1 Network Architecture and Design, E.1 Data Structures, H.2.8 Database Applications, I.2.9 Robotics, I.2.11 Distributed Artificial Intelligence, K.4.1 Public Policy Issues, K.4.2 Social Issues, K.6.4 System Management
Keywords and phrases amorphous computing, decentralized spatial computing, distributed algorithms, location privacy, organic computing, self-organization, sensor/actuator networks, situation awareness, smart materials, spatial analysis
Digital Object Identifier 10.4230/DagRep.3.12.17

1 Executive Summary

Matt Duckham
Stefan Dulman
Jörg-Rüdiger Sack
Monika Sester

The aims of Dagstuhl Seminar 13492, "Geosensor Networks: Bridging Algorithms and Applications," were to advance research into, and application of geosensor networks by enhancing interdisciplinary and cross-domain collaboration. The premise of the Seminar was that the potential for useful, practical applications of geosensor networks (wireless sensor networks tasked with monitoring changes in geographic space) are being held back by the enormous diversity of applications and expertise connected with different facets of geosensor networks. The result today is many niche solutions to specific problems, where what is needed are a few general solutions to broader problems.

More specifically, the diversity of concepts, approaches, and tools used in connection with geosensor networks is inhibiting more rapid and fruitful research and development. Examples of this diversity include:

- *Ontologies, representations, and models*: The diversity of uses of geographic information leads inexorably to diversity in the different ontologies, representations, and models commonly used to conceptualize that information.
- *Algorithms and data structures*: The plethora of models, algorithms, data structures, and architectures that exist in the literature are frequently incompatible, founded on divergent assumptions and inconsistent approaches.
- *Applications*: While the potential applications of geosensor networks are legion, today we lack agreement on a set of applications that together encompass and illustrate the bulk of issues faced by all applications.
- *Benchmarks, tools, and technologies*: Perhaps more than any other single issue, a lack of consensus on core benchmark data sets, problems, simulation systems, and software tools inhibits convergence in research and application.

Participants

To begin to address this diversity, and bridge the gap between theory and application, the Seminar participants represented a broad spectrum of disciplines, including computer science, geographic information science, computational geometry, statistics, artificial intelligence, pervasive computing. The seminar had strong groundings in previous Dagstuhl Seminars, including Seminars 10491 and 12512 on Representation, Analysis and Visualization of Moving Objects, and Seminar 06361 on Computing Media and Languages for Space-Oriented Computation. However, we were very pleased that the Seminar also attracted a significant proportion of newer Dagstuhl attendees: more than half the attendees had attended at most one Dagstuhl seminar before, with around one quarter of Seminar participants attending their first ever Dagstuhl.

Bringing together this diversity of backgrounds, expertise, and experiences was central to the core aim of building bridges between related fields, and was central to the success of the Seminar.

Format

The seminar was structured around three complementary perspectives: models and algorithms; benchmarking and applications; and teaching and curricula. The objective of the models and algorithms perspective was to survey, catalog, and compare the ontologies, representations, algorithms, and data structures that are fundamental to computing in geosensor networks. Through the benchmarks and application perspective, the seminar aimed to improve comparability and compatibility in models and algorithms, as well as connect existing models and algorithms more directly to practical uses. Focal applications included emergency response, intelligent transportation, smart materials, and environmental monitoring. As a capstone, the teaching and curricula perspective aims to distill and collate the collected expertise in models, algorithms, benchmarks, and applications into a coherent body of knowledge: a “library” of core concepts and techniques for computation with and application of geosensor networks.

The seminar focused less on presenting individual lectures than on achieving its objectives above through collaborative discussions and activities. The organizers invited three speakers with diverse backgrounds to give longer talks (40 minutes) and lead the subsequent discussions. The three speakers were René Doursat (CNRS, Paris, discussing organic computing), Thomas Kirste (Universität Rostock, discussing situational awareness and intention recognition using sensed data), and Edzer Pebesma (Universität Münster, on spatial data analysis with sensor data). As well as providing an introduction to the breadth and depth of ideas related to the field, the speakers were able inspire the participants and spark many subsequent discussions.

The majority of the seminar then focused on workshop-style discussion and break-out groups. In this way the seminar aimed to elicit answers to the question of what are the essential elements of computing with geosensor network. The aim was be to advance the field through consensus on priorities as well as providing opportunities for new innovations to emerge from new collaborations. The working groups' discussions and conclusions are summarized in this report (Section 5). Broadly, the working groups' focuses included spatial computing (e.g., self-organization and smart materials); applications of sensor networks (e.g., developing countries and big data); social issues (e.g., privacy); education (e.g., teaching resources and curriculum); and data and benchmarking.

However, even though the primary focus was on discussion and collaboration, the program still allowed time for short focus talks from participants (up to 10 minutes for senior researchers, or up to 15 minutes for junior researchers—researchers were able to self-select as to whether they regarded themselves as junior or senior). A summary of the focus talks given by participants is also contained in this report. All the speakers were asked to address one of the three Seminar perspectives (models and algorithms; benchmarking and applications; and teaching and curricula), as can be seen from section 4.

Outcomes

The participants were highly satisfied with the quality of the seminar. Many and diverse research results were presented during the Seminar, surveyed in the following sections. As with many Dagstuhl Seminars, the new collaborations and results of those collaborations are ongoing. However, amongst the key findings and ongoing collaborations, we highlight:

- Considerable progress has been made in recent years in the areas of (decentralized) spatial computing. This includes advances in the bottom-up design of distributed and decentralized algorithms. However, in contrast top-down aggregate programming techniques offer an important advantage over more conventional decentralized programming, in that they are substantially less complex for developers to use. Whatever future advances in this area may hold, the Seminar participants were agreed in that decentralization is a means, but never an ends. A focus on the behavior of a distributed geosensor network as a whole, rather than the rules required to generate that behavior, should always be the focus. In this respect, the focus on emergent behaviors found in spatial computing would seem ideally suited.
- One particular area of progress at the seminar was in teaching and curriculum. Despite the wide range of expertise and academic backgrounds of the participants, the Seminar exposed the considerable commonality and agreement around the fundamental concepts behind geosensor networks. This convergence was evident in the recent development of tools (such as the Proto aggregate programming language) and the publication of text books on the subject of geosensor networks.

- In contrast, one area of particular difficulty was in benchmarking and data sets. The availability of data sets is unquestionably increasing, as evidenced by several different data sets that were made available at, and through the preparatory work by participants in advance of the Seminar. However, the wide diversity of requirements for data sets across different applications continues to defy standardization or convergence on a small set of benchmarks. Issues such as validation and ground truthing; requirements for massive data sets with thousands or millions of sensors; metadata and provenance; and privacy issues were all various inhibitors to the development of a small set of benchmark data sets.
- Finally, the Seminar highlighted the numerous practical, social, and environmental challenges that still remain to truly bridging the gap between theoretical algorithms and practical applications, such as cost and deployment strategies, privacy, and environmental pollution. Although these issues remain largely unsolved longer-term research problems, they are already the explicit focus of several new collaborations that have directly resulted from this Seminar.

2 Table of Contents

3 Invited Talks

3.1 Architectured self-organized systems: Toward the best of both worlds by "morphogenetic engineering"

René Doursat (CNRS, FR)

Engineering is torn between an attitude of strong design and dreams of autonomous devices. It wants full mastery of its artifacts, but also wishes these artifacts were much more adaptive or "intelligent." Meanwhile, the escalation in system size and complexity has rendered the tradition of rigid top-down planning and implementation in every detail unsustainable. In this context, natural complex systems (large sets of elements interacting locally and behaving collectively) can constitute a powerful source of inspiration and help create a new generation of "self-x" systems, these properties being mostly absent from classical engineering.

This talk showed other avenues of bio-inspired design stressing the importance and benefits of a genuine *self-organization in architectured systems*——as exemplified by the growth of multicellular organisms or the nests of social insects. A new field of research was presented, "morphogenetic engineering" (ME; http://doursat.free.fr/morpheng.html), which explores the artificial design of complex morphologies that can reproducibly arise without centralized or externalized control. Potential applications range from swarm robotics to distributed software, techno-social networks and geosensor networks to synthetic biology. What they have in common is a myriad of hardware/software/bioware/geoware agents that can be programmed to dynamically build structures on the sole basis of peer-to-peer communication and local computation.

3.2 Symbolic behavior models for intention recognition in situation-aware assistance

Thomas Kirste (Universität Rostock, DE)

One objective of situation aware assistance is to support persons that need to perform structured activities by providing guidance and automation. In general, this requires the estimation of the actual activity sequence and the underlying intentions, from noisy and ambiguous sensor data. Unfortunately, obtaining the training data required for building probabilistic inference systems that employ sequential state estimation is often very expensive. Therefore, it is of interest to use prior knowledge on the causal structure of action sequences for building a system model. Symbolic causal models are one option for doing this, for instance based on the STRIPS paradigm used in the domain of symbolic planning. This talk outlined the basic mechanisms for attaching a probabilistic semantics to such a symbolic model and discussed the applicability of this approach to activity recognition problems of realistic complexity. The results obtained on the reconstruction of everyday activities from wearable accelerometer data (specifically, an instrumental activity of daily living) suggest that such problems can indeed be successfully tackled by symbolic modeling methods.

3.3 Sensor networks from a data analytic perspective

Edzer Pebesma (University of Münster, DE)

This talk surveyed a number of distinct practical use cases over the past 20 years connected with sensor networks and analysis of the data they produce.

In the first use case [1] we analyze data from groundwater quality monitoring networks, which are sampled yearly to assess the availability and development of long-term quality water resources. We tried to map spatial variability of groundwater quality and its dynamics, and found that for several components (e.g. nitrate) either the noise was too high, or the Nyquist frequency was not met and predictions were highly inaccurate, even when predicting average concentrations for 4km $\times$ 4km blocks. By only presenting results in terms of 95% prediction intervals, we avoided the misinterpretation of highly uncertain predicted values.

In the second use case [2], the goal was to evaluate, optimize or design from scratch a monitoring network that should detect nuclear outbreaks from various possible origins in an as short as possible time, on a country-wise or region-based basis. This led to a spatial design problem, and by representing space as a regular grid on a fine-enough resolution, to a discrete optimization problem. The implicit random field to be sampled was generated by deterministic plume models, covering a relevant range of sources and weather conditions.

In the third use case [3], we evaluated rural background air quality data collected at a European scale, and assessed whether for the interpolation of yearly aggregated values one should model and interpolate the daily values directly, or rather aggregate station time series to yearly aggregates and interpolate those.

The fourth use case addresses ambient variables and car properties (e.g. temperature, fuel consumption) measured from cars, and collected on a data sharing site called envirocar.org[1]. Understanding data from mobile sensors, interpreting and analyzing them poses new problems, not addressed in mainstream spatial statistics literature. Analyzing such data also leads to the problem which spatial predictions (interpolations) and aggregations are meaningful, and which are not [4]: we can interpolate car engine temperatures, but doing this for a location and time where there is no car is meaningless. Averaging over several cars may also be not meaningful.

Finally, the talk addressed how collaborative software development takes place in the R project [5], where sharing software and analysis scripts has become the rule, rather than the exception. This opens up solutions to the problem of reproducing scientific results, as far as they concern computational aspects, as well as benchmarking computational procedures.

References

1 Pebesma, E.J. and J.W. de Kwaadsteniet. Mapping groundwater quality in the Netherlands. *Journal of Hydrology*, 200:364–386, 1997.

2 Helle, K.B. and E.J. Pebesma. Stationary sampling designs based on plume simulations. In Jorge Mateu and Werner G. Müller (eds.) *Spatio-temporal Design: Advances in Efficient Data Acquisition*, Chapter 14, 2012.

3 Gräler, B., L. Gerharz, and E.J. Pebesma. Spatio-temporal analysis and interpolation of PM10 measurements in Europe. ETC/ACM Technical Paper 2011/10; Released: 2012/01/30.

[1] See http://www.envirocar.org/ and http://rpubs.com/edzer/envirocar2

4 Stasch, C., S. Scheider, E.J. Pebesma, and W. Kuhn. Meaningful spatial prediction and aggregation. *Environmental Modelling & Software*, 51:149—165, 2014.
5 Bivand, R.S., E.J. Pebesma, and V. Goméz-Rubio *Applied Spatial Data Analysis with R*, Second edition. Springer, NY, 2013.

4 Focus Talks

Theme: Models and Algorithms

4.1 Data mining techniques in sensor networks: Summarization, interpolation, and surveillance

Annalisa Appice (University of Bari, IT)

Joint work of Annalisa Appice, Anna Ciampi; Fabio Fumarola; Donato Malerba
Main reference A. Ciampi, D. Malerba, F. Fumarola, A. Appice, " Data Mining Techniques in Sensor Networks: Summarization, Interpolation and Surveillance," SpringerBriefs in Computer Science, 105p., ISBN 978-1-4471-5454-9, Springer, 2014.
URL http://www.springer.com/computer/database+management+%26+information+retrieval/book/978-1-4471-5453-2

Sensor networks consist of distributed devices, which monitor an environment by collecting data (light, temperature, humidity, ...). Each node in a sensor network can be imagined as a small computer, equipped with the basic capacity to sense, process and act. Sensors act in dynamic environments, often under adverse conditions.

Typical applications of sensor networks include monitoring, tracking, and controlling. Some of the specific applications are photovoltaic plant controlling, habitat monitoring, traffic monitoring, ecological surveillance. In these applications, a sensor network is scattered in a (possibly large) region where it is meant to collect data through its sensor nodes.

While the technical problems associated with sensor networks have reached a certain stability, managing sensor data brings numerous computational challenges in the context of data collection, storage, and mining. In particular, learning from data produced from a sensor network poses several issues: sensors are distributed; they produce a continuous flow of data, eventually at high speeds; they act in dynamic, time-changing environments; the number of sensors can be very large and dynamic. These issues require the design of efficient techniques for processing data produced by sensor networks. These algorithms need to be executed in one step of the data, since typically it is not always possible to store the entire data set, because of storage and other constraints.

The focus of this talk was to provide an idea of data mining techniques in sensor networks. We have taken special care to illustrate the impact of data mining in several network applications by addressing common problems, such as data summarization, interpolation, and surveillance. We propose a clustering technique to summarize data and permit the storage and querying of this amount of data, produced by a sensor network in a server with limited memory. Clustering is performed by accounting for both spatial and temporal information of sensor data. This permits the appropriate trade-off between size and accuracy of summarized data. Data are processed in windows. Trend clusters are discovered as a summary of each window. They are clusters of georeferenced data, which vary according to a similar trend along the time horizon of the window [2]. We can use trend clusters to interpolate missing data. The estimation phase is performed by using the *inverse distance weighting approach* [1]. Finally, we can also use trend cluster for fault detection in photovoltaic plants [3].

References

1 A. Appice, A. Ciampi, D. Malerba, and P. Guccione. Using trend clusters for spatiotemporal interpolation of missing data in a sensor network. *Journal of Spatial Information Science*, 6(1):119–153, 2013.

2 A. Appice, A. Ciampi, and D. Malerba. Summarizing numeric spatial data streams by trend cluster discovery. *Data Mining and Knowledge Discovery* 2013, in press.

3 A. Ciampi, A. Appice, D. Malerba, and A. Muolo. An intelligent system for real time fault detection in PV plants. *Smart Innovation, Systems and Technologies*, 12:235–244, 2012.

4.2 Programming distributed algorithms using computational fields

Jacob Beal (BBN Technologies, Cambridge MA)

A geosensor network can typically be viewed as a spatial computer: a collection of devices distributed through space such that the difficulty of moving information between devices is strongly dependent on the distance between them, and the "functional goals" of the system are generally defined in terms of the system's spatial structure.

It is often useful to use decentralized approaches in implementing such networks, particularly when network infrastructure is unavailable or unreliable. Pragmatically, however, the difficulty of implementing complex and reliable distributed algorithms has formed a major barrier to using such techniques.

Aggregate programming techniques can lower this barrier, by allowing engineers to program a geosensor network as though it were centralized, and then automatically compile that program into an adaptive distributed algorithm implementing the equivalent.

We have found much success in using a continuous space abstraction, the amorphous medium, which views a geosensor network as a discrete sampling of the continuous space that is being monitored. Field calculus provides a space-time universal basis set for aggregate programming under this model, which enables a number of forms of implicit safe composition, adaptivity, and scalability. We have elaborated this model into an implementation as the Proto programming language (http://proto.bbn.com), and applied this to programming a broad range of platforms, including sensor networks, swarm robotics, agent-based simulations, and engineered biological cells.

4.3 Smart materials: From sensor networks to computational meta-materials

Nikolaus Correll (University of Colorado, Boulder CO)

Joint work of Apice, Annalisa; Beal, Jacob; Correll, Nikolaus; Dulman, Stefan; Doursat, René; Papp, Zoltan; Sack, Jörg-Rüdiger; Sester, Monika

Materials can become truly "smart" by combining properties of the carrier polymer with sensing, actuation, and computational, networked logic. Depending on the complexity of the underlying distributed algorithms, such materials can be purely reactive like the human skin, implement distributed processing and actuation like the retina or colon, respectively, or even become conscious like the human brain. Research in smart materials and their application will be driven by the desire for multi-functionality of existing materials and surfaces. For example, a vinyl floor can easily be extended by pressure sensors and colored lights, enabling applications such as emergency routing, indoor navigation, occupancy mapping or playful activities using comparably simple, gradient-based distributed algorithms. Here, the additional cost for sensing, actuation, and computation might become marginal given the resulting multi-functionality. Similarly, enhancing fabrics by sensing, actuation and computation might enable a new generation of garments that can not only monitor, but affect posture and physiological parameters of their wearer. Other functional materials might have the ability to shape-change, monitor structural health and self-repair, or even assemble and disassemble, allowing repurposing of an object's functionality and its enclosing space. Scalability and robustness requirements of such materials strongly motivate the use of fully distributed, self-organizing algorithms, many of which have begun to be studied by the swarm robotics, sensor network, and computational geometry communities. In addition to challenges in material science engineering and manufacturing, truly smart materials will require composing the existing suite of distributed algorithms into aggregates of increasing complexity that run on larger numbers of computational nodes with minimalist capabilities.

4.4 Organic building evacuation support system: A vision

Sabrina Merkel (KIT: Karlsruhe Institute of Technology, DE)

A rapidly growing world population and increasingly dense settlements demand ever-larger and more complex buildings from architects and engineers. However, current evacuation support systems in these buildings are comparatively outdated. Their analogous and static character makes evacuation support easily overlooked in panic and prevents route guidance from being adaptive to the current evacuation situation, such as recognizing and avoiding congestions. The Organic Building Evacuation Support System (OBESS) introduced in this talk is a concept for solving these issues by using mobile devices to guide people to safe exits. The observer/controller architecture from the research area of organic computing was introduced. It was shown how this architecture can be applied to mobile devices in order to achieve adaptive and self-organized localization and evacuation path planning that is robust and controllable at the same time. Furthermore, OBESS provides for a central control unit, which is meant to perform building-specific optimization in order to adapt the system to the building's individual requirements. One example for such an optimization approach is the improvement of the positions of anchor nodes in the building, which are used for localization of the mobile devices. In this talk, an evolutionary algorithm developed for this purpose was described. Moreover, four fitness evaluation criteria are introduced which are tested in experiments. The presented research shows that optimizing the coverage of anchor nodes is not enough to yield good localization results and that the localization algorithm, as well as the applied performance metric, has crucial impact on optimal anchor placements. Apart from the presented research, a variety of research topics that arise in this area are proposed in this talk, which have to be addressed in order to achieve a long overdue changeover in evacuation support of modern buildings.

Theme: Applications and Benchmarking

4.5 Geosensor networks applications

Steve Liang

This session discussed the potential for a "killer app" in geosensor networks. The session started by identifying two real-world applications:

1. monitoring wild mammal populations and migration in Northern Alberta [1] (more specifically, monitoring woodland caribous and wolves); and
2. pipeline safety monitoring.

Both applications present similar challenges, including the remote sites, power constraints, extreme weather, and potentially high maintenance costs, etc.

Domain knowledge plays critical roles for each both applications. For example, domain knowledge is required to determine the parameters to be observed and the installation location of the sensor nodes. Deploying large-scale geosensor networks in real-world remains a challenging task.

Other potential killer applications might include "smart space garbage": how might decentralized geosensor networks be used to track debris and detritus in the space?

References

1 Corey J.A. Effects of petroleum exploration on woodland caribou in northeastern Alberta. *The Journal of Wildlife Management*, 61(4):1127–1133, 1997.

4.6 Runtime reconfiguration in sensor/actuator networks

Zoltan Papp (TNO, The Hague, NL)

In order to address today's societal challenges (e.g. environmental management, intelligent mobility, safety/security, sustainable urban living, etc.) monitoring and control of large-scale, complex, dynamical systems are fundamental technologies. These monitoring and control systems realize sophisticated signal processing algorithms (state estimation, pattern recognition, optimal control, decision making, etc.) in distributed configurations and many times are deployed in difficult to access, hostile environments. Under these circumstances failures, changes in operational conditions and changes in user needs are not exceptions but should be considered as nominal operation. Runtime adaptation capabilities (covering self-organization, self-optimization) have to be realized to extend the operational envelope of the system.

The design of runtime adaptable systems poses new challenges both for the signal processing and for the system architecture aspects. The presentation surveyed the approaches for automatic runtime adaptation and identifies the design and implementation challenges. A model-based methodology was introduced, which greatly simplifies the development of these complex artifacts and thus enables the successful introduction of the runtime adaptation concept in various of application domains. The distinguishing features of the targeted systems are serious resource constraints (e.g., processing capabilities, communication, energy supply) and the presence of demanding non-functional requirements, such as timeliness, robustness, lifetime and, the capability of handling system "evolution." The "practical aspects" the runtime reconfiguration were also addressed (representation, reasoning and developer support). The talk concluded with identifying a number of open issues, such as resource aware and distributed algorithms for optimization, decision making and control, measurement of "costs" and "usefulness," guaranteeing pre-defined closed-loop behavior.

4.7 Live geography: Fusing data from technical and human sensors

Bernd Resch (University of Heidelberg, DE)

The predicted rise of geo-sensor webs has not taken place as quickly as estimated, which prevents many research efforts from being carried out due to lacking real-time base data. Additionally, we are currently witnessing the rapid emergence of user-generated data in a variety of geo-social networks, the mobile phone network, or micro-blogs. This human-generated data can potentially complement sensor measurements to a significant degree.

In terms of user-centered sensing approaches, three concepts can be distinguished:

1. "People as sensors" defines a measurement model, in which measurements are not only taken by calibrated hardware sensors, but in which also humans can contribute their subjective "measurements," such as their individual sensations, current perceptions, or personal observations.
2. "Collective sensing" tries not to exploit a single person's measurements and data, but analyzes aggregated anonymized data coming from collective sources, such as Twitter, Flickr, or the mobile phone network.
3. "Citizen science" stands for a human-based approach to science where citizens contribute semi-expert knowledge to specific research topics.

Particular challenges in fusing data from human and technical sensors include standardization (on data, service and method levels), data assimilation (resolution, aggregation, etc.), the combination of methods from geoinformatics and computer linguistics (to extract information from user-generated data), quality assurance (both for technical and human sensors), and the consideration of privacy issues (data ownership, storage, optimum aggregation levels, etc.).

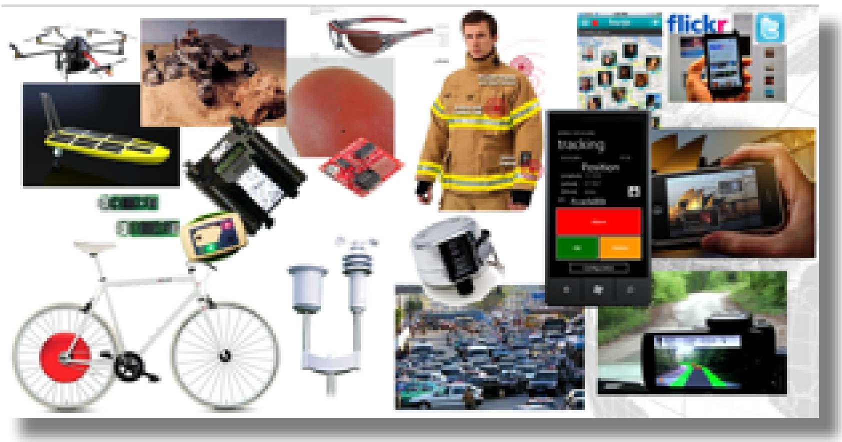

Figure 1 Live geography: Diverse sources of sensor data.

References

1 Resch, B. People as sensors and collective sensing—contextual observations complementing geo-sensor network measurements. In J. Krisp *Advances in Location-Based Services*, Springer, Berlin. pp. 391–406, 2013.

2 Resch, B., A. Zipf, P. Breuss-Schneeweis, E. Beinat, and M. Boher. Towards the live city—paving the way to real-time urbanism. *International Journal on Advances in Intelligent Systems*, 5(3&4): 470–482, 2012.

4.8 Environmental testbeds for dependable sensor networks

Kay Römer (TU Graz, AT)

Sensor networks provide a substrate to realize applications in several domains of utmost importance for our society, including surveillance of critical infrastructures, smart cities, smart grids, and smart healthcare. However, many of these applications are only possible if sensor networks provide dependable performance. Application-specific guarantees on network performance parameters such as data delivery reliability and latency must be given for all system operation conditions. Failure to meet these requirements at all times may lead to reduced user satisfaction, increased costs, or to critical system failures. Unfortunately, existing sensor network technologies mostly follow a best-effort approach and do not offer guaranteed performance.

The major hurdle to providing dependable sensor is that their operation is deeply affected by their surrounding environment. Environmental properties such as electromagnetic (EM) radiation, ambient temperature, and humidity have significant impact on achievable network performance. Sensor network communication has to deal with interference from other communication networks such as WiFi, Bluetooth, RFID and from other systems such as microwave ovens or engine ignition systems. The resulting interference may lead to message loss, which in turn leads to an increase in latency and energy consumption. Environmental temperature also deeply affects the operation of WSAN. We have shown in previous work that temperature variations in a deployment may lead to failing transmissions during hot periods. Not only are these environmental conditions hard to predict for a given deployment site, they also may largely vary from one deployment site to another, thus hindering scalable deployment of IoT applications as every new deployment site requires costly customization. For example, climatic conditions and the use of frequency bands varies heavily across cities and countries.

In order to design sensor networks that can provide certain performance guarantees despite changing environmental conditions, there is a need for testbeds with realistic environmental effects, where protocols and applications can be run on real sensor network hardware under repeatable and realistic environmental conditions. In this talk, two such testbeds were presented—TempLab and JamLab—where user-defined temperature and interference conditions can be created. In TempLab, sensor nodes are equipped with infrared heating lamps that can be controlled via wireless dimmers to create an accurate temperature profile that varies over time and space. In JamLab, a subset of the sensor nodes in the testbed is used to record and playback interference patterns with high accuracy without the need for additional hardware.

Theme: Teaching and Curricula

4.9 Education materials: Lessons learned and opportunities

Susanne Bleisch (The University of Melbourne, AU)

GITTA (Geographic Information Technology Training Alliance, http://www.gitta.info) was initially funded by the Swiss Virtual Campus Initiative to create comprehensive and in-depth online teaching materials in geographic information science and technology. While the original eLearning hype pushed the hopes of saving money it was soon realized that this is difficult to achieve and many of the early eLearning projects disappeared after funding run out. GITTA was able to survive by opening up its content under a creative commons license and transferring from the original project consortium to a supporting association that welcomes new partners. Additionally, the project employs the eLesson Markup Language eLML (http://www.elml.ch) to achieve Learning Management System independence and output format flexibility. The XML-based language eLML supports content creation by offering the ECLASS structure. From these eLearning projects was learned that the discussion between the different project partners about learning objectives and lesson contents is very valuable. However, each teacher also has his or her own style and does not necessarily like to use predefined teaching materials. Additionally, it is often difficult to fit a series of longer learning units such as lessons into the available time slots or the lesson levels or contents do not correspond with the students' prior knowledge. For the usage of the case studies it is crucial that the employed data sets are available.

From a user survey conducted in 2010 it is known that the GITTA materials are actively used for teaching and learning. However, it is also known that quite often the materials are used as a source of inspiration or single illustrations or animations from the GITTA materials are used in lectures. Based on this knowledge it seems natural to propose that the focus of creating educational materials in the area of geosensor networks should take on the form of small and flexible learning objects or resources. Learning objects are more complete but small units that are ideally structured (for example according to the eLMLs ECLASS principle), aligned with learning objectives and described through metadata (for example the IMS Learning Resource Meta-data Specification, http://www.imsglobal.org/metadata/). Even smaller are learning resources, which may consist of single illustrations, simulations, questions, etc. To improve discoverability, learning resources should ideally also be aligned with learning objectives and described through metadata. A repository of learning objects or resources allows using them flexible in a range of settings such as face-to-face lectures, flipped classroom or potentially mobile learning. Each user has the freedom to choose and to incorporate into his or her teaching what seems to fit with fewer difficulties regarding teaching style, level, or time allocations.

4.10 Education and decentralized spatial computing

Matt Duckham (The University of Melbourne, AU)

Main reference M. Duckham, "Decentralized Spatial Computing: Foundations of Geosensor Networks," 320p., ISBN 978-3-642-30852-9, Springer, 2013.
URL http://link.springer.com/book/10.1007/978-3-642-30853-6
URL http://ambientspatial.net/book

Learning about decentralized spatial computing requires the development of specific domain skills and knowledge across a number of important areas, including computing and distributed systems; spatial information and analysis; and algorithm design and analysis. In addition to these specialized topics, decentralized spatial computing links well to several general skills for students of computer science, engineering, and geographic information science. Specifically, teaching the design and testing of decentralized spatial algorithms can support the development transferable skills including:

1. logical thinking and analysis, requiring careful and structured problem solving and system engineering, in the design of the algorithms;
2. constructive critical thinking and self-criticism, necessary in the adversarial analysis of algorithms; and
3. validating conclusions through the construction of objective and statistical experimental evidence to support the properties of an algorithm.

The information on http://ambientspatial.net/book to complement the book [1] contains many examples of algorithm and practical code to support teaching and learning.

References

1 Matt Duckham *Decentralized Spatial Computing: Foundations of Geosensor Networks*. Springer, Berlin. 2013.

5 Working Groups

5.1 Smart Materials

Stefan Dulman, Jacob Beal, Rene Doursat, Kay Romer, and Nikolaus Correll

The breakout session started by identifying examples of smart materials and listing related work examples, including:

- strong objects (tables);
- biological materials (synthetic cells; bacteria produced materials);
- intelligent fabric;
- smart facades;
- room dividers; and
- health related things, prosthetics, smart chair, etc.

Recent funding programs have showcased many more examples, such as DARPA WarriorWeb (e.g., smart suits with accelerometers and sensors; smart skin; competitors to exoskeleton; shirts with microphones, that locate where sound comes from; and applications to geriatrics, which identify when some someone is falling or losing their balance).

This led to a refinement of the definition of smart materials from the perspective of our community. In short, features of smart materials include:

- autonomous components;
- physical coupling of devices;
- morphological computation;
- a combination of the analog and digital computing;
- a focus on "programmable";
- distinct from existing approaches like body area networks by sheer numbers, amorphous characteristics, imprecision, and price.

In turn, some of the problems facing smart materials include: What can the engineering provide (properties of materials)? What can be manufactured with what cost? and What can we program (programming emergence , a focus in this workshop).

As a result of the brainstorming, the group explored a specific example of smart materials, arriving at the concept of a programmable fabric equipped with a set of sensors (stress and shear, temperature, distance, medical, audio, physiology sensors); actuators (vibration, stiffness, torque, force, heat, light, biochemical, thermal); and outputs (pose information = somatosensory computing, pose support, mecano-transduction).

Applications for this programmable fabric include:

- mechanical actuation, e.g., posture manipulation;
- monitoring of sleeping, e.g., detecting scratching and preventing it.
- posture correction;
- sport training, e.g., teaching/observing/correcting moves;
- fall prevention, e.g., detect loss of equilibrium, geriatric support;
- mechanical actuation, e.g., tactile signaling;
- haptic feedback and virtual reality;
- feeling i-materials, e.g., signal strengths for access point deployment or radiation detection.

Finally, in order to create a prototype for the programmable fabric, the following building blocks would be needed on top if the material science engineering:

- distributed operating systems (viral programming, maintenance, debugging, error detection);
- several programming layers (start from the aggregates and end up with a medic's user interface);
- a library of adaptive computations (needed to raise the level of programming abstraction, e.g., doctors should be able to program);
- modeling in simulations (connect to tools used in material science, such as abaqus);
- centralized control (centralized calibration of sensors, self-assessment, safety and certification); and
- identifying who is in control (e.g., shirt manufacturer, doctor, user).

5.2 Self-organization and geosensor networks

Rene Doursat, Stefan Dulman, Jacob Beal, Kay Romer, Zoltan Papp, Sabrina Merkel, and Nikolaus Correll

Questions posed in this group included: What makes the topic of self-organization interesting? How does self-organization fit the field of geosensor networks? Would it be possible to import ideas and tools from other domains?

Self-organization can be defined as a collection of objects that form and maintain a structure. Temporal efficiency needs to be part of the definition. Clear requirements are needed before deciding that a distributed approach is to be preferred. Typically, there exist trade-offs between functionality,spatial constraints, and resources (communication).

Real systems need teal-time predictability. Combining control systems with emergence must lead to a system collectively predictable. The decentralized features of a practical system need to "be hidden under the hood." Users are not interested in the centralized/decentralized system design approach. Instead, users care about efficiency, reliability, robustness, and modularity.

For example, wilderness monitoring is an interesting application for distributed systems (wild fires, precision agriculture, ocean/environment monitoring). A second application is field infrastructure (emergency and disaster monitoring, and the developing word). Smart materials is a third application area, bringing in the question of high density nodes. The switch from physical to wireless would be a loss of information density. These scenarios involve dynamic/volatile architectures (nodes/links come and go) but one wants to see a stable functionality and have clear guarantees.

However, it was agreed that self-organization should be used only if there is a clear incentive to do so. The effort must not outweigh the benefits.

5.3 Benchmarking and data challenge

Matt Duckham, Annalisa Appice, Edzer Pebesma, and Monika Sester

The development of benchmark data sets and an associated data challenge in the area of geosensor networks would serve two purposes:
1. Increasing comparability and reproducibility of research results, allowing many researchers to test algorithms and analysis on the same data set.
2. Helping with outreach and enlarging the community, showcasing relevant problems.

Geosensor network data is typically low-accuracy, high-resolution data. In the benchmarking and data challenge session, we identified certain key characteristics of an idea benchmark data set and associated data challenge:
- simple accessible data formats (e.g., csv files, database tables);
- raw, unprocessed sensor readings;
- real in preference to simulated data;
- data about both static and mobile sensor platforms;
- context descriptions and ground truth information (e.g., events that are known to have occurred in the region/period covered by the data).

A data challenge would then need to identify and make accessible suitable data, develop applications and questions for the challenge, and help provide a forum for collaboration and communication of results (for example via a workshop, a challenge associated with a major conference, or a Dagstuhl seminar).

One of the planned activities following on from this seminar is the development of a data challenge. For example, http://www.marinetraffic.com/ makes accessible shipping AIS data which might be collated for a geosensor network data challenge, such as:
- based on historical movement, forecast vessels in a specific region or port in the future;
- detect anomalies, such as suspicious movements or AIS spoofing.

5.4 GSN4D: Geosensor networks for development: Applications in smart grid and e-health

Hedda Schmidtke, Edzer Pebesma, Daniel Fitzner, Thomas Kirste, Sabrina Merkel, and Martijn Warnier

Geosensor networks (GSN) have a number of applications for facilitating international development programs. International development is a rewarding field of application that contributes to saving human lives and secure international stability and security. This breakout session focused on two particular application areas: smart grid and e-health. We discussed how standard application scenarios for the developed world match to scenarios for development. In the smart grid area, standard GSN application scenarios include:
- How can GSN be used to provide energy from renewable sources during peak times and control energy flow? and
- How can end-users be made aware of their energy usage, so as to incentivize desirable behavior of consumers?

In a developing world scenario, similar research questions are relevant. But here, avoiding peaks in demand is crucial to avoid power outages and damage to equipment. Inclusion of renewable energy sources in this scenario facilitates economic development and increased life quality for citizens. Many of the practical challenges that hinder innovation in energy systems in the developed world are less of a concern in developing countries. Utilities in developing countries may be more open towards integrating smart energy systems, as these systems may contribute to improved reliability of services and protection of assets. Also, end-consumers in the developing world may be more likely to wish to understand how they can have a positive influence on the operation of the grid.

There is a wide area of opportunities for GSN research in this application scenario. Communication protocols and distributed processing can be studied well with a smart grid application scenario. Sensor nodes can be attached to power infrastructure for monitoring and controlling the energy flow. Moreover, the energy grid can provide power to sensors alleviating constraints arising from battery life concerns in other geosensor network applications. Ad-hoc networks and self-organizing network structures are a focus of GSN research that is particularly interesting in a developing world scenario. But mobile communications protocols, such as GSM, are also viable alternatives for implementing communication for GSN research with a focus outside of the communications area. All research scenarios in the area of GSN for smart grid applications, such as decentralized processing of information, and GSN for monitoring, error prediction, and localization of errors carry over to a developing world scenario. Some differences exist in the economic perspective on smart grids, since consumers in the developing world are in general more willing to adopt the new technologies without economic incentives, as a more reliable grid is an incentive per se.

However, consumers in both scenarios need to be made aware of energy consumption and grid operations with easy-to-use interfaces. Another question is how smart appliances or smart plugs can be economically sensible options in a developing world scenario. Interesting options would be to integrate smart plug functionality into plugs for spike protection, or to integrate smart meters into pre-paid meters regularly installed into households.

Among the algorithms and models used in the domain are market simulation and prediction models and algorithms, e.g. multi-agent systems and game theory, as well as network models to simulate the grid structure as a spatially embedded network, and to model the energy flow. Special distributed processing simulations and algorithms are needed to test smart grid GSNs, integrating the wide range of interdisciplinary perspectives on the problem: from the perspective of the electrical engineer to the economist and social science perspective.

A similarly exciting and worthwhile application domain in the larger area of applications for international development is e-health. GSN in this domain arise as crowd sensing or human sensing networks. A large proportion of the population in developing countries owns a mobile phone. Successful applications for the case of Rwanda, e.g., are the RapidSMS implementation for maternal and child health[2] and the mUbuzima (Kinyarwanda for m-health) project for monitoring parameters for the Milennium Development Goals[2] set by the UN. RapidSMS is used for monitoring pregnancies in remote villages and also for intervention in case of complications. In case of an emergency, the system can provide advice and automatically requests an ambulance to bring the mother to the nearest hospital. Like in the case of smart grid applications, applications in e-health require an interdisciplinary approach with GSN research integrating medical and/or pharmaceutical perspectives.

[2] http://www.unmillenniumproject.org/goals/

References

1 Krogh, B.H. and H.R. Schmidtke. Electrical power in Africa: Challenges and opportunities. *IEEE Smart Grid Newsletter*, 12 2012.

2 Ngabo, F., J. Nguimfack, F. Nwaigwe, C. Mugeni, D. Muhoza, D.R. Wilson, J. Kalach, R. Gakuba, C. Karema, and A. Binagwaho. Designing and implementing an innovative SMS-based alert system (RapidSMS-MCH) to monitor pregnancy and reduce maternal and child deaths in Rwanda. *The Pan African medical journal*, 13, 2012.

5.5 Data integration and big data from sensors

Annalisa Appice, Susanne Bleisch, Allison Kealy, Thomas Kirste, Steve Liang, Hedda R. Schmidtke, and Monika Sester

Geosensor networks (GSN) produce large amounts of data. In order to make data usable a number of questions have to be addressed. Two key challenges are scalability and integration of heterogeneous data. An important goal is the discovery of hidden knowledge in data through queries, browsing, or exploration of data. Technologies to access this hidden knowledge are machine learning and logical reasoning methodologies or combinations of these. To consider an example, assume a suspicious boat has been reported: how can a GSN be used to determine what has happened? In this case information may be available from sensors deployed in the water, remote sensing, boat AIS providing ID, position, and velocity, plus material from social networking platforms, such as twitter and Flickr. Spatial and temporal indexing play a key role to limit the search space, assuming that locality implies causal relation between events.

5.6 Privacy issues for geosensor networks

Bernd Resch, Edzer Pebesma, Matt Duckham, Jörg-Rüdiger Sack, and Martijn Warnier

Geosensor networks collect data. Typically a lot of different data is collected, some of which might be privacy-sensitive because it can be linked back to individuals. For this reason engineers working in the field of geosensor networks should be aware of the privacy issues involved.

Ultimately, when designing geosensor networks, the privacy of individuals related with the network should be respected. It is not always clear who's privacy should be protected, because data can be linked back to individuals in many ways, including combining geosensor data with other networks. Several issues are of importance here:

- End users should be able to balance privacy needs with performance measures, i.e., by offering more detailed information about oneself, the offered service might perform better.
- Decentralized architectures inherently have the potential to allow more fine grained control/choice for end-users about the type of information they share and the granularity of that information. However, decentralized solutions do not guarantee more privacy by default.

- Several forms of data obfuscation can help the privacy of end users. For example, inaccurate time stamps in data might help privacy of individuals. This comes at the cost of less reliable—and probably overall less useful—data.
- Privacy awareness of end-users remains an issue, though applications such as Snapchat (http://snapchat.com) might raise the privacy awareness of users.
- Privacy transparency needs further exploration. For most applications it is not clear how they process data (and who they share it with), making most applications opaque from a privacy perspective.

We concluded that guaranteeing privacy is a complex and difficult issue, with no general solutions for the all cases, but with some possibilities for application- or context-specific solutions.

5.7 Teaching geosensor networks

Warnier, Martijn

This group discussed some of the issues related to teaching geo sensor networks. Besides the (obvious) conclusion that a book such as [1] is a nice starting point for this (especially with the integration of a programming language like NetLogo), showing students how to think "from the sensor" and stresses decentralized design, the group also identified the importance of several other issues. These include showing (and letting students play with) real hardware. Playing with hardware should be stimulated at the beginning of an (introductory) course, so that students become familiar with issues such as energy constraints and get a feeling for networks. For example, it might be possible to let students design their own routing protocols, which are then peer-reviewed by the students and then implemented in NetLogo. This stimulates them to think for themselves and let them appreciate the complexities involved. A competitive element might also stimulate students (e.g., Who can build the best sensor network?). In turn, this might require a teaching platform and/or real data.

References

1 Matt Duckham *Decentralized Spatial Computing: Foundations of Geosensor Networks*. Springer, Berlin. 2013.

6 Open Problems

Matt Duckham, Jacob Beal, René Doursat, Thomas Kirste, Steve Liang, Zoltan Papp, Edzer Pebesma, Hedda Schmidtke, and Monika Sester

Finally, an open problems session yielded the following range of open problems:

JB Can the Nyqvist Theorem be generalized to universal space-time computation?
ZP Can we guarantee that we find a solution for self-organization within a particular time constraint?

EP How do I generalize dense attributed trajectories? Cartographic generalization operates only on the spatial projection, not space-time.
TK Is it possible to define SVD (single value decomposition) for high dimensional categorical state space using a generative definition?
JB What API will make it easy for sensor network engineers to use aggregate programming?
SL How do I discover the sensors fit for my application and search results need to be fast, fresh, complete, relevant?
MS Is there a space-time search engine (i.e., a GIR engine that also accounts for time)?
RD/HS Don't we have to give up on (some of) our needs for assurance/validation/feasibility of implementation in exchange for qualitative leap in functionality/benefits strong self organization? And if yes, why aren't we doing it? This may alternatively be decomposed into two questions. What is the best trade off between certainty and predictability in engineering self-organizing algorithms? and 2. How to use a machine for brainstorming, computer aided creativity?

7 Future Visions

René Doursat

This closing session invited all participants to describe their "dreams" or long-term motivations in their research on decentralized sensor networks. What did they imagine (or would they like to see) happening in several years or decades from now? What are to them the ultimate endeavors of this cross-disciplinary field? Additionally, what are the potential dangers or abuses that could also emerge from these new technologies?

The answers were extremely diverse, illustrating the wide scope of application domains and techniques. A useful classification that emerged from this session relies on the spatial scale of the systems:

System scales (from smallest to largest)

- in the body (< 1mm): "swarms" of nanotech particles (medication delivery, targeted cures);
- on the body (1–10cm): sensor/actuator "skin" or "suit," wearable computing;
- on furniture, on a wall (1m): "touch-table," life-size videoconferencing;
- across a home (5–10m): assisted living, accident detection, ambient intelligence;
- in a building (10–100m): automated regulation of utilities, emergency evacuation;
- in a crowd: opportunistic networking via human/cellphone carriers;
- across a city (1–10km): traffic control and optimization (lights, traffic jams, detours);
- across a social network: Wiki-like mobile collaboration;
- across a region, territory (10–100km): forest fires, agriculture, ecology of fauna/flora;
- in a population: disaster preparedness, epidemic, emergency response; and
- in the ocean or atmosphere (100–1000km): currents, climate, environment, pollution.

A few of these potential applications were discussed further. Here follows a summary, which attempts to categorize them into major domains and topics (with many crossovers, naturally):

Body, health, and assisted living

- automated health assessment by navigating through the environment, for example, detecting a cold leads to medication or a recommendation to visit the pharmacy.
- smart clothes combining sensing and actuating, for example, a sore muscle or headache is detected leading to a proposed massage; or a baby's respiration or a senior's gait/balance are monitored.
- automated physical therapy, for example, pose recognition and posture manipulation; or rehabilitation, and regaining patterns of motion.
- body implants with sensors, for example "smart" hip bone prosthesis.
- enabling a better connection between senior people and their entourage, for example alerting a medical center in case of emergency or simply being able to say hello to family displayed on the wall.

Human crowds, social networks, populations

The main idea in this category is to create an "opportunistic" peer-to-peer network that relies on the human carriers, essentially via their PDAs: cell phones, smart watches, or some other specific device (such as a bracelet given out at a concert).

- *entertainment*: a crowd at an event (sports, concert, demonstration), for example, PDAs emit colored light signals—seen from above, the whole crowd creates waves, patterns or entire videos on a giant "amorphous" display.
- *security*: law enforcement, emergency response, for example, counting how many people there are or guiding the crowd to an evacuation route, via PDAs showing them the way.
- *information, knowledge*: Wiki-like collaboration MANETS (mobile ad-hoc networks), for example, the empowerment of anyone to innovate and contribute to project (in a "prosumer" fashion), such as "mobile journalism."
- *disaster management*: information to support responses to earthquakes, floods, nuclear accidents, epidemics, etc.

The main challenges here include the needs for interoperability between people's contributions; and dealing with privacy problems (and/or adapting to changes in expectations and approaches to privacy).

Smart materials, mechanical/civil engineering, aeronautics

Many of the visions centered around new materials and smart paints or coatings. A myriad of immobile tiny sensor-actuator particles, densely strewn on a surface (stuck in place), collectively performing "amorphous computing" formed the infrastructure for these visions, including:

- at home, for example a sensitive carpet, touch-walls, touch-tables, etc.
- engineering structures, such as self-monitoring bridges and aircraft fuselages.
- beyond assessing, actuating and optimizing a structure, for example modifying an aircraft wings' profile during flight ("smart/soft aeronautics").
- a myriad of mobile tiny sensor/actuator particles ("smart dust"), which either passively glide along with the winds or currents, or are self-propelled and can steer themselves, for example in environmental applications (forests, rivers, crops, etc.).

For the latter applications in particular, there was the potential risk of unwanted "nano-pollution." Creating a degradable paint or dust sensor particles might avoid this accumulating pollution problem. Degradation might happen spontaneously, or upon a global termination signal, using abiotic (organic chemistry, biomimicry) or biotic (leveraging the self-organizing power of bacteria, synthetic biology) means.

Theoretical and technical challenges

Altogether, what is really needed is not just smart "sensing," but a whole sensing-computing-actuating cycle in each element.

A major question raised was whether there was any advantage in using a decentralized algorithm as opposed to a centralized algorithm. The question was posed: "Which applications warrant decentralization?" Conversely, which applications can be implemented in a classical way, i.e., by a computing method that has complete knowledge of all the data available in the system?

- *Answer 1*: applications where space plays an important role in restricting the flow of information, so that there are spatial constraints to the movement of information.
- *Answer 2*: where there is a need to optimize resources (primarily energy), even if their could technically be a central computing device, i.e., we want to do more with less resources.

It was also argued that, although important, this question actually should not concern the programmer. Decentralization, if any, should be "transparent" or "under the hood." Ideally, there would be a "compiler" able to transform any central algorithm into a set of rules distributed over a myriad of agents. Going even further, such systems should allow anyone, not just skilled programmers, to innovate by using and transforming the sensing information. Users should not only be consumers, but also "natural" programmers who could instantly create new applications.

Next, was the lower-level issue of how decentralized algorithms could be implemented. How do elements collaborate in a peer-to-peer fashion and make decisions collectively? Typically, spatial structures, such as gradients of "hop counters," are an important fundamental mechanism of self-organization.

Finally, a methodological and epistemological question was raised. With all these future network computing systems, is it not the case that traditional engineering requirements for assurance of results, validation/proof/checks, and feasibility of implementation may need to be relaxed in favor of a step-change in functionality provided by genuine self-organization? In other words, it may be that "organic"-like systems, where function is the emergence of a great collective of small elements, will never be fully predictable or controllable. Then, it is not a matter of guaranteeing a service, but rather setting an acceptable level of probability of success. We must also be prepared for that.

Participants

- Annalisa Appice
University of Bari, IT
- Jacob Beal
BBN Tech. – Cambridge, US
- Susanne Bleisch
The University of Melbourne, AU
- Nikolaus Correll
Univ. of Colorado – Boulder, US
- René Doursat
CNRS – Paris, FR
- Matt Duckham
The University of Melbourne, AU
- Stefan Dulman
TU Delft, NL
- Daniel Fitzner
Leibniz Univ. Hannover, DE
- Allison Kealy
The University of Melbourne, AU
- Thomas Kirste
Universität Rostock, DE
- Steve Liang
University of Calgary, CA
- Marco Mamei
University of Modena, IT
- Sabrina Merkel
KIT – Karlsruhe Institute of Technology, DE
- Zoltan Papp
TNO – The Hague, NL
- Edzer Pebesma
Universität Münster, DE
- Bernd Resch
Universität Heidelberg, DE
- Kay Römer
TU Graz, AT
- Jörg-Rüdiger Sack
Carleton Univ. – Ottawa, CA
- Hedda R. Schmidtke
Carnegie Mellon University – Moffet Field, US
- Monika Sester
Leibniz Univ. Hannover, DE
- Martijn Warnier
TU Delft, NL