09371 Abstracts Collection
Algorithmic Methods for Distributed Cooperative Systems
— Dagstuhl Seminar —

Sándor Fekete[1], Stefan Fischer[2], Martin Riedmiller[3] and Suri Subhash[4]

[1] TU Braunschweig, D
s.fekete@tu-bs.de
[2] Universität Lübeck, D
fischer@itm.uni-luebeck.de
[3] Universität Osnabrück, D
riedmiller@informatik.uni-freiburg.de
[4] Univ. California - Santa Barbara, USA
suri@cs.ucsb.edu

Abstract. From 06.09.09 to 11.09.09 the Dagstuhl Seminar 09371 "Algorithmic Methods for Distributed Cooperative Systems" was held in Schloss Dagstuhl – Leibniz Center for Informatics. The purpose of this workshop was to bring together researchers from different disciplines whose central interest is in both algorithmic foundations and application scenarios of distributed cooperative systems. 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. Algorithms, cooperative systems, sensor networks, multi-robot systems, applications.

1 Executive Summary – Algorithmic Methods for Distributed Cooperative Systems

A standard scientific method for *understanding* complicated situations is to analyze them in a top-down, hierarchical manner. This approach also works well for *organizing* a large variety of structures; that is why a similar hierarchical, centralized approach has worked extremely well for employing computers in so many aspects of our life: Data is gathered, processed, and the result is administered by one central authority.

On the other hand, the structures in our modern world are getting increasingly complex. The resulting challenges have become so demanding that it is impossible to ignore that a large variety of systems are based on a very different principle: The stability and effectiveness of our modern political, social and economic structures relies on the fact that they are based on decentralized, distributed and self-organizing mechanisms. This paradigm shift is also reflected in a variety of modern computing systems, which work in a distributed manner, based on local (and thus: incomplete) information and interaction, and implement the results in a localized fashion; as opposed to a variety of social or economic situations, we may assume that the individual components of such a system are not primarily selfish, but pursue a joint goal that is to be reached in collaboration.

The purpose of this workshop was to bring together researchers from different disciplines whose central interest is in both algorithmic foundations and application scenarios of distributed cooperative systems. In particular, participants from the following communities were present:

AF **Algorithmic Foundations.** When developing a systematic method for solving an algorithmic problem by a cooperating set of loosely coupled processors, the result is a distributed algorithm. One of the resulting consequences is incomplete information, for which an increasing number of algorithmic aspects have been studied. Moreover, a number of additional issues are considered, such as communication complexity, timing issues, and the amount and type of information that is available to individual processors.

SN **Sensor networks.** In recent time, the study of wireless sensor networks (WSN) has become a rapidly developing research area, both from the theoretical and the practical side. Typical scenarios involve a large swarm of small and inexpensive processor nodes, each with limited computing and communication resources, that are distributed in some geometric region; communication is performed by wireless radio with limited range. From an algorithmic point of view, these characteristics imply absence of a central control unit, limited capabilities of nodes, and limited communication between nodes. This requires developing new algorithmic ideas that combine methods of distributed computing and network protocols with traditional centralized network algorithms. In other words: How can we use a limited amount of strictly local information in order to achieve distributed knowledge of global network properties? Just now, an important set of additional challenges for sensor networks is starting to emerge from mobile nodes, making it necessary to deal with additional problems arising from network dynamics.

RT **Multi-robot systems.** Multi-robot systems consist of several individual robots, either identical or heterogeneous. Among the scenarios for teams or swarms of autonomous robots are robot soccer (RoboCup), rescue missions, exploration and other complex tasks that can be carried out in a distributed fashion. Beyond the technical aspects of perception, behaviors, learning, and action, the most interesting issue in the context of this interdisciplinary seminar are various modeling aspects that are getting quite close to those faced in areas such as SN: after all, a sensor-equipped robot becomes quite similar to a mobile sensor node, and both face similar difficulties, but also possibilities.

AP **Application scenarios.** In order to provide further scenarios for challenges and discussions, we included a selection of researchers from other application areas; among these was

- Traffic: making use of car-to-car communication, it has become possible to provide online, up-to-date local information and coordination. What challenges can be tackled by making use of these possibilities?
- Biology: swarm behavior of animals has developed over millions of years. What lessons can be learned from such behavior?

These four aspects were subdivided into algorithmic foundations (provided by AF), two specific areas (SN and RT) that form the link between pure theory and real-life applications, and a variety of real-life challenges (provided by AP) that can serve as goals and benchmarks for the other scientific work.

Quite naturally, there was some amount of overlap between these four areas in terms of individual researchers, as various scientists combine theory and practice, to a varying degree. Despite of the previous distinction between the different fields, a variety of aspects implied that similar problems were faced, so that a dialogue between the researchers turned out to be quite fruitful.

Each of the fields both benefitted and contributed:

AF Theoretical methods of distributed algorithms and algorithms with incomplete information form the basis for the algorithmic work on the application scenarios (AP), as well as the problems arising in both SN and RT. Fundamental insights and results turn out to be useful for the development of practical methods, but also show basic obstacles for obtainable results. Conversely, application scenarios help to focus the theoretical algorithmic work, and lead to the identification of new kinds of problems.

SN The practical side of sensor networks gives rise to a number of quite specific scenarios for distributed algorithms. Many of these problems consider stationary nodes, for which a variety of aspects enjoy an increasing amount of understanding; however, there is growing demand for dealing with mobile sensor nodes (in particular when dealing with scenarios from the application areas AP), requiring extensions of theoretical work (AF), but also leading to a growing similarity with scenarios faced by RT.

RT The ever-improving technology and control for autonomous robot systems has become more and more sophisticated, and advanced to the point where there is an increasing demand for higher-level, algorithmic methods (AF). The aspects of dynamics, communication and outside information give rise to a number of quite challenging scenarios (AP); moreover, problems like the exploration of unknown territory by a swarm of robots still require a lot of algorithmic work (AF). This is where the similarity to systems of mobile nodes in a sensor networks (SN) are striking; it is obvious that the exchange between both communities was quite beneficial, as some of the fundamental challenges are surprisingly similar.

AP The application areas described above provide both a collection of grand challenges and a reality check for the theoretical methods by AF and the specific methods developed by SN and RT; on the other hand, solutions and insights by AF, SN, and RT gave rise to completely new possibilities for mastering those challenges.

The workshop brought together 35 researchers from nine countries. The 20 presentations, varying in length, covered a large variety of topics. Owing to the combination of different research areas, there were a number of survey talks and discussion session, but also a variety of individual research presentations. In addition, there was sufficient time for informal discussion and small-scale interaction.

Overall, participants agreed that the interdisciplinary nature of the workshop was quite fruitful. There was strong interest in repeating this event with a similar combination of fields and researchers in the not-too-distant future.

From Hierarchical Reactive Control to Hierarchical Planning for Humanoid Soccer Robots

Sven Behnke (Universität Bonn, DE)

In my talk, I introduce the RoboCup Humanoid League and our team NimbRo.

In particular, I cover the hierarchical reactive behavior control architecture that we developed for complex dynamic tasks, such as playing soccer. In this architecture, behaviors of increasing complexity are executed on the higher layers, which run on slower timescales.

Our team won several of the RoboCup competitions, including KidSize 2007 and 2008, and TeenSize 2007 and 2009.

While the abstraction used in the hierarchical reactive control scheme has many benefits, it also prevents the full utilization of the available hardware, e.g. when approaching a ball for kicking without slowing down. I motivate why planning a constant number of actions on each of these layers can improve system performance in such situations. On the example of local multiresolution path planning, I give first results of such a hierarchical planning approach.

Keywords: Hierarchical reactive control, hierarchical planning, robot soccer, RoboCup

On the Fairness of Probabilistic Schedulers for Population Protocols

Ioannis Chatzigiannakis (RACTI, Greece, GR)

We propose a novel, generic definition of *probabilistic schedulers* for population protocols. We design two new schedulers, the *State Scheduler* and the *Transition Function Scheduler*. Both possess the significant capability of being *protocol-aware*, i.e. they can assign transition probabilities based on information concerning the underlying protocol. We prove that the proposed schedulers, and also the *Random Scheduler* that was defined by Angluin et al. are all fair with probability 1. We also define and study *equivalence* between schedulers w.r.t. *performance* and *correctness* and prove that there exist fair probabilistic schedulers that are not equivalent w.r.t. to performance and others that are not equivalent w.r.t. correctness. We implement our schedulers using a new tool for simulating population protocols and evaluate their performance from the viewpoint of experimental analysis and verification. We study three representative protocols to verify stability, and compare the experimental time to convergence with the known complexity bounds. We run our experiments from very small to extremely large populations (of up to 10^8 agents). We get very promising results both of theoretical and practical interest.

Keywords: Population Protocols, Fairness, Probabilistic Schedulers, Communicating Automata, Sensor Networks, Experimental Evaluation

Joint work of: Chatzigiannakis, Ioannis; Dolev, Shlomi; Fekete, Sándor; Michail, Othon; Spirakis, Paul

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2428

Improving the Flow of Traffic

Sándor Fekete (TU Braunschweig, DE)

Traffic is one of the most important complex systems of our modern world, with several levels of complexity reaching from individual actions of a driver, over local phenomena like density fluctuations and traffic jams, regional and temporal traffic patterns, all the way up to long-range traffic development and regulation.

In recent years, tremendous progress has been made in understanding the dynamics of traffic flow and traffic congestion; however, a number of serious obstacles still prevent efficient coordination and regulation of traffic. Three of the most serious impediments have been the incompleteness of input data, the computational intractability of forecasting the behavior of real-life traffic consisting of huge numbers of vehicles, and the lack of local communication (and thus: cooperation) between drivers.

With the advances of modern communication technologies, it has become possible to keep track of virtually all data of driving vehicles. Understanding traffic as a complex system that is based on local interaction suggests studying distributed computing approaches for controlling traffic phenomena. Finally, wireless ad-hoc networks allow real-time interaction and data exchange between adjacent vehicles.

In this talk I describe how these tools can be used to improve the flow of congested traffic. Simulations suggest that this may allow saving about 40% of energy for cars in a traffic jam.

Keywords: traffic flow, wireless communication, distributed cooperation, energy consumption.

Joint work of: Fekete, Sándor; Hendriks, Björn; Schmidt, Christiane; Tessars, Chris; Fischer, Stefan; Hellbrück, Horst; Wegener, Axel

Cooperative Multi-Agent Systems from the Reinforcement Learning Perspective – Challenges, Algorithms, and an Application

Thomas Gabel (Universität Freiburg, DE)

Reinforcement Learning has established as a framework that allows an autonomous agent for automatically acquiring – in a trial and error-based manner – a behavior policy based on a specification of the desired behavior of the system.

In a multi-agent system, however, the decentralization of the control and observation of the system among independent agents has a significant impact on learning and its complexity.

In this survey talk, we briefly review the foundations of single-agent reinforcement learning, point to the merits and challenges when applied in a multi-agent setting, and illustrate its potential in the context of an application from the field of manufacturing control and scheduling.

Keywords: Multi-agent reinforcement learning, decentralized control, job-shop scheduling

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2010/2426

Interaction networks and colony organization in ants

Deborah M. Gordon (Stanford University, US)

An ant colony operates without central control. Each ant uses only local information, mostly odor, and no ant can make global assessments about what needs to be done. Through the local decisions of individuals, colonies accomplish task allocation, adjusting the numbers of ants devoted to each task to the current situation. An ant decides where to go and what to do based on its recent experience of brief interactions with other ants. Most interactions consist of antennal contact, in which one ant smells the other. In an interaction, one ant perceives the chemical signature of another ant's task, but does not give or receive any instructions. The rate of interaction, rather than any information transferred, influences task decisions. Interaction networks explain how seed-eating ants organize where they forage and how they regulate the intensity of foraging according to the availability of food. Whether an ant leaves the nest on its next foraging trip depends on the rate at which ants return with food. Since most of a forager's trip is spent searching for food, and a forager does not return until it finds food, the rate of forager return is a measure of how quickly food can be found. Current research investigates the algorithm used to link forager return rate to foraging activity in these harvester ants. Task allocation in harvester ants depends on colony age and size; older, larger colonies show more stable dynamics than younger, smaller ones. This raises the question of how colony size affects

interaction networks and thus colony behavior. Ants are a huge group of more than 11,000 species with diverse ecology and behavior. Other research in our group investigates how interaction networks function in ant species that collect food by continually modifying an established highway system rather than by individual searching.

Keywords: Interaction network

Efficient Self Stabilizing Data Structures

Riko Jacob (TU München, DE)

Topological self-stabilization is an important concept to build robust open distributed systems (such as peer-to-peer systems) where nodes can organize themselves into meaningful network topologies.

The goal is to devise distributed algorithms that converge quickly to such a desirable topology, independently of the initial network state. This paper proposes a new model to study the parallel convergence time. Our model sheds light on the achievable parallelism by avoiding bottlenecks of existing models that can yield a distorted picture. As a case study, we consider local graph linearization—i.e., how to build a sorted list of the nodes of a connected graph in a distributed and self-stabilizing manner. We propose two variants of a simple algorithm, and provide an extensive formal analysis of their worst-case and best-case parallel time complexities, as well as their performance under a greedy selection of the actions to be executed.

Joint work of: Gall, Dominik; Jacob, Riko; Scheideler, Christian; Schmid, Stefan; Täubig, Hanjo; Richa, Andrea

Local Strategies for Connecting Stations by Small Robotic Networks

Barbara Kempkes (Universität Paderborn, DE)

Consider a group of m stations with fixed positions in the plane and a group of n mobile robots, called relays, aiming at building a communication network between the stations consisting of as few relays as possible. We present a strategy for dimensionless, identical (anonymous), oblivious and disoriented relays with limited viewing radius for constructing such a network. A relay does not communicate with others, its decision—whether to remove itself from the system, or where to move—consists only of the relative positions of its neighbors within its viewing radius. We provide a theoretical analysis of worst-case scenarios and upper and lower bounds for the number of relays used by the strategies.

Mapping and Exploration for Search and Rescue with Humans and Mobile Robots

Alexander Kleiner (Universität Freiburg, DE)

Urban Search And Rescue (USAR) is a time critical task since all survivors have to be rescued within the first "golden" 72 hours. One goal in Rescue Robotics is to support emergency response by mixed-initiative teams consisting of humans and robots. Their task is to explore the disaster area rapidly while reporting victim locations and hazardous areas to a central station that plans for rescue missions.

To fulfill this task efficiently, humans and robots have to map disaster areas jointly while coordinating their search at the same time. Basically, they have to solve autonomously in real-time the problem of Simultaneous Localization and Mapping (SLAM), consisting of a continuous tracking problem, and a discrete data association problem. However, in disaster areas these problems are extraordinarily challenging due to rough terrain, making tracking difficult, as well as due to the unstructured environment, complicating the recognition of visited places. SLAM methods work with the principle of map improvement through loop closure, i.e. to improve the map globally each time places have been re-observed, which turns out to be suboptimal during exploration. Furthermore, coordination and mapping has to function reliably even if communication links are cut off.

Following the vision of combined multi-robot and multi-human teamwork, core problems, such as position tracking on rough terrain, mapping by mixed teams, and decentralized team coordination with limited radio communication, are directly addressed by this talk. More specific, I will introduce RFID-SLAM, a method for robust and efficient loop closure in large-scale environments that utilizes RFID technology for data association.The method is capable of jointly improving multiple maps from humans and robots in a centralized and decentralized manner without requiring team members to perform loops on their routes. Thereby positions of humans are tracked by PDR (Pedestrian Dead Reckoning), and robot positions by odometry, respectively. The joint-graph emerging from these trajectories serves as an input for an iterative map optimization procedure. The introduced map representation is further utilized for solving the centralized and decentralized coordination of large rescue teams. On the one hand, a deliberate method for combined task assignment and multi-agent path planning, and on the other hand, a local search method using the memory of RFIDs for coordination, are proposed.

Methods introduced in this talk were extensively evaluated in outdoor environments and official USAR testing arenas designed by the National Institute of Standards and Technology (NIST). Furthermore, some were an integral part of systems that won multiple awards at international competitions, such as the RoboCup world championships.

Keywords: SLAM, Coordination, Exploration, RFID, USAR, Rescue, Mapping, Mixed-Initiative Teams

Full Paper:
http://www.informatik.uni-freiburg.de/~kleiner/papers/kleiner_diss08.pdf

See also: A. Kleiner, *Mapping and Exploration for Search and Rescue with Humans and Mobile Robots*, Ph.D. thesis, University of Freiburg, 2008

Distributed Vision with Smart Pixels

Alexander Kroeller and Christiane Schmidt (TU Braunschweig, DE)

We study a problem related to computer vision: How can a field of sensors compute higher-level properties of observed objects deterministically in sublinear time, without accessing a central authority? This issue is not only important for real-time processing of images, but lies at the very heart of understanding how a brain may be able to function.

In particular, we consider a quadratic field of n "smart pixels" on a video chip that observe a B/W image. Each pixel can exchange low-level information with its immediate neighbors. We show that it is possible to compute the centers of gravity along with a principal component analysis of all connected components of the black grid graph in time $O(\sqrt{n})$, by developing appropriate distributed protocols that are modeled after sweepline methods.

Our method is not only interesting from a philosophical and theoretical point of view, it is also useful for actual applications for controling a robot arm that has to seize objects on a moving belt. We describe details of an implementation on an FPGA; the code has also been turned into a hardware design for an application-specific integrated circuit (ASIC).

Keywords: distributed vision, distributed algorithms, sublinear algorithms, principal component analysis, sweepline algorithms.

Joint work of: Fekete, Sándor; Kroeller, Alexander; Schmidt, Christiane; Fey, Dietmar; Komann, Marcus; Reichenbach, Marc

Extended Abstract:
http://portal.acm.org/citation.cfm?doid=1542362.1542410

The Role of Centralized Brokers in Distributed Computations

Alejandro Lopez-Ortiz (University of Waterloo, CA)

We describe three examples of rather successful, widely deployed distributed systems, namely the World Wide Web, the Border Gateway Protocol (BGP) and cluster computing farms (such as Amazon S3). In the case of the web, usability was greatly improved by the addition of search engines (Lycos, OpenText, and later Google, Bing) which perform the role of centralized brokers within the web architecture. In the case of BGP companies such as Internap—which use a hybrid distributed+centralized approach—can easily outperform fully distributed BGP-only solutions. Lastly Amazon S3 has been most successful with small companies with highly variable loads (or unpredictable loads), rather than with companies with massive but stable CPU requirements. This matches the analysis by A. L.-O. and Stiffelman [2001] as well as, independently by Jim Gray.

This highlights the need for a fresh approach to distributed applications using a judicious combination of distributed and centralized brokers, with emphasis on applications which are novel and cannot easily be implemented in-house either on a sequential computer or using a local cluster.

Keywords: Distributed computing, WWW, routing, cloud computing

Information-Driven Sensor Network Architecture versus Traditional Layered Node Architecture

Ingrid Moerman (Ghent University, BE)

Wireless sensor networks are used for increasingly complex applications, such as wireless building automation, medical monitoring, tracing of persons and assets, etc. These applications demand more and more functionalities from the underlying sensor network. They are often deployed on strongly heterogeneous nodes, and require adaptive and reliable end-to-end services. Such requirements cannot be supported at a protocol level, but should be part of the overall sensor network architecture. However, no architecture currently exists that supports heterogeneity, easy protocol-integration and QoS as part of its architectural design.

We recently have proposed an alternative sensor node architecture, called information-driven architecture (or IDRA). This architecture is based on the notion that protocols should be simplified to their two main tasks: exchanging information and interacting with the relayed information. By intelligently manipulating this information in a distributed fashion, the system can support advanced network requirements for next-generation sensor network applications.

The information driven approach has the following key advantages:

- by using a packet facade for packet interactions, protocol logic is decoupled from packet representation;
- rather than statically wiring protocols, protocols are dynamically selected (based on protocol-provided filters);
- by providing a shared, system-wide packet queue, the overall memory footprint is reduced and system-wide
- QoS can be enforced;
- heterogeneity is promoted since protocols can be added to a node according to its capabilities;
- and finally, by efficiently combining the information exchanges, the number of transmitted packets can be strongly reduced.

Keywords: Wireless sensor networks, network architecture

Joint work of: Moerman, Ingrid; De Poorter, Eli

Multiple Random Walks and Interacting Particle Systems

Tomasz Radzik (King's College London, GB)

We study properties of multiple random walks on a graph under various assumptions of interaction between the particles. To give precise results, we make our analysis for random regular graphs. For k independent walks on a random regular graph G, the cover time is asymptotic to $C(G)/k$, where $C(G)$ is the cover time of a single walk. If the walks can communicate when meeting at a vertex, we show that, for most starting positions, the expected time for k walks to broadcast a single piece of information to each other is asymptotic to $[(r-1)/(r-2)] * [(2 \ln k)/k] * n$, where r is the degree of the nodes and n is the number of nodes. The asymptotics of the time needed for gossiping is not known. We also consider properties of walks where there are two types of particles, predator and prey, or where particles interact when they meet at a vertex by coalescing, or by annihilating each other.

Full Paper:
http://www.springerlink.com/content/p700274230357002/

Joint work of: Cooper, Colin; Frieze, Alan; Radzik, Tomasz

Multi-Agent Learning in Team of Robots

Martin Riedmiller (Universität Freiburg, DE)

In this talk I'll provide an overview of our research on multi-robot coordination and cooperation by the means of Reinforcement Learning. In particular, I'll discuss that along the line of research in our RoboCup robotic soccer team Brainstormers. Both in simulation league and in Middlesize league, machine learning concepts were integrated into classical software systems with the explicit goal to increase the competitveness of our agents. Within 10 years of active participation, we won 6 titles, several 2nd and 3rd places in annual worldchampionships. The efficiency of our algorithms has drastically improved since then. Now, even learning directly on real robots is possible.

Keywords: Multi-agent cooperation, cooperative reinforcement learning

Debugging Sensor Networks in the Field

Kay Roemer Uwe (Universität Lübeck, DE)

Sensor networks are deeply embedded into the physical world in order to monitor real-world phenomena. The resulting close interaction between the real world and an embedded sensor network entails many challenges.

In particular, the dynamics of a real-world environment negatively affects sensor nodes (e.g., physical strain) and radio communication between them, often leading to failure of the sensor network to meet application goals. As these real-world aspects are hard to model in the lab, failures often manifest only after deployment in the field. Novel distributed debugging techniques and tools are required to monitor the performance of a sensor network in situ on the deployment site, to identify and remove the causes of insufficient performance. As sensor networks run on a very constrained resource and energy budget, these techniques should only consume a small fraction of this limited budget. Through this small "keyhole" of resources (network bandwidth, energy), the state of the sensor network has to be observed and controlled in a minimally invasive manner. This talk discusses the challenges of minimally invasive in-field debugging of sensor networks and presents approaches developed by the author, namely passive inspection and passive distributed assertions.

Keywords: Sensor networks, debugging, distributed assertions, sniffing

The Fundamental Limits of Broadcasting in Wireless Mobile Networks

Paolo Santi (CNR - Pisa, IT)

In this talk, we investigate the fundamental properties of broadcasting in mobile wireless networks. In particular, we characterize broadcast capacity and latency of a mobile network, subject to the condition that the stationary node spatial distribution generated by the mobility model is uniform. We first study the intrinsic properties of broadcasting, and present a broadcasting scheme, called RippleCast that simultaneously achieves asymptotically optimal broadcast capacity and latency, subject to a weak upper bound on the maximum node velocity. This study intendedly ignores the burden related to the selection of broadcast relay nodes within the mobile network, and shows that optimal broadcasting in mobile networks is, in principle, possible. We then investigate the broadcasting problem when the relay selection burden is taken into account, and present a combined distributed leader election and broadcasting scheme achieving a broadcast capacity and latency which is within a $\Theta((\log n)^{1+\frac{2}{\alpha}})$ factor from optimal, where n is the number of mobile nodes and $\alpha > 2$ is the path loss exponent. However, this result holds only under the assumption that the upper bound on node velocity converges to zero (although with a very slow, poly-logarithmic rate) as n grows to infinity.

To the best of our knowledge, our is the first paper investigating the effects of node mobility on the fundamental properties of broadcasting, and showing that, while optimal broadcasting in a mobile network is in principle possible, the coordination efforts related to the selection of broadcast relay nodes lead to sub-optimal broadcasting performance.

Keywords: Wireless networks; mobile networks; broadcast capacity; broadcast latency; SINR interference model

Joint work of: Resta, Giovanni; Santi, Paolo

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2427

Stabilizing Consensus with the Power of Two Choices

Christian Scheideler (University of Paderborn, DE)

Consensus problems occur in many contexts and have therefore been intensively studied in the past. In the standard consensus problem there are n processes with possibly different input values and the goal is to eventually reach a point at which all processes commit to exactly one of these values. We are studying a slight variant of the consensus problem called the stabilizing consensus problem. In this problem, we do not require that each process commits to a final value at some point, but that eventually they arrive at a common value without necessarily being aware of that. This should work irrespective of the states in which the processes are starting. Coming up with a self-stabilizing rule is easy without adversarial involvement, but we allow some T-bounded adversary to manipulate any T processes at any time. In this situation, a perfect consensus is impossible to reach, so we only require that there is a time point t and value v so that at any point after t, all but up to $O(T)$ processes agree on v, which we call an almost stable consensus. As we will demonstrate, there is a surprisingly simple rule for the standard message passing model that just needs $O(\log n \log \log n)$ time for any $\sqrt{n}$-bounded adversary and just $O(\log n)$ time without adversarial involvement, with high probability, to reach an (almost) stable consensus from any initial state. A stable consensus is reached, with high probability, in the absence of adversarial involvement.

Keywords: Distributed consensus

Joint work of: Doerr, Benjamin; Goldberg, Leslie Ann; Minder, Lorenz; Sauerwald, Thomas; Scheideler, Christian

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2429

Minimalist Mapping and Monitoring

Suri Subhash (Univ. California - Santa Barbara, US)

We propose and investigate a minimalist sensing model for autonomous group of simple robots, and prove possibility and impossibility results for basic exploration tasks such as mapping and monitoring of a physical environment.

Keywords: Minimalism, visibility, mapping, art gallery theorem

Joint work of: Subhash, Suri; Peter, Widmayer; Elias, Vicari; Matus, Mihalis

See also: Simple Robots with Minimal Sensing: From Local Visibility to Global Geometry, International Journal of Robotics Research, Vol 27, Sept. 2008.

Sensor Networks meet Organic Computing

Dirk Timmermann (Universität Rostock, DE)

The talk introduces sensor networks and their close relationship to typical social and biological networks. Hence, self-x properties and strategies can be found in and newly applied to sensor networks. As an example of the benefits of organic principles for organic networks, cluster algorithms and scale free alike networks are presented.

Distributed Algorithms

Roger Wattenhofer (ETH Zürich, CH)

An introduction to distributed algorithms, in particular local algorithms. Essentially a practice talk of my SSS 2009 invited talk.

Keywords: Local Algorithms, Self-Stabilization, Lower Bounds, Upper Bounds, MIS

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2425

Full Paper:
http://disco.ethz.ch/publications/sss09_LSW.pdf

Local Actions and Global Structure of Complex Networks

Katharina A. Zweig (Universität Heidelberg, DE)

In this talk we examine how the environment of a complex network can influence the local actions of single agents such that a globally optimal network structure arises. We analyze two scenarios: one in which agents try to achieve a more central position over time, and one in which agents try to recover from the loss of neighbors. In the first scenario we prove that feeding back the maximal distance to an agent can lead to an exponential runtime until all agents are satisfied with their position. In contrast, feeding back the sum of distances to each agent reduces the runtime to a polynomial function. In the second scenario, we prove that there exists a universal, local, and oblivious repair mechanism that results in a robust network structure after removal of nodes. Both scenarios show that a self-organized network needs to be analyzed with respect to its environment, and we expect new exciting questions to arise from a joint analysis of networks and their environment.

Keywords: Network structure, real-world networks, network analysis

Full Paper:
http://www-pr.informatik.uni-tuebingen.de/mitarbeiter/katharinazweig/downloads/299-lehmann.pdf

Full Paper:
http://www.ninasnet.de/Publications/PDFs/Zweig-WandererBetweenWorlds.pdf

See also: K.A. Lehmann, Michael Kaufmann, Evolutionary Algorithms for the Self-Organized Evolution of Networks, Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'05), 2005, K.A. Zweig and Karin Zimmermann, Wanderer between the Worlds – Self-Organized Network Stability in Attack and Random Failure Scenarios, Proceedings of the 2nd Second IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO'08), 2008