Report from Dagstuhl Seminar 14051

Algorithms for Wireless Communication

Edited by

Guy Even[1], Magnús M. Halldórsson[2], Yvonne Anne Pignolet[3], and Christian Scheideler[4]

1 **Tel Aviv University, IL,** `guy@eng.tau.ac.il`
2 **Reykjavik University, IS,** `mmh@ru.is`
3 **ABB Corporate Research, Baden, CH,** `yvonne-anne.pignolet@ch.abb.com`
4 **Universität Paderborn, DE,** `scheideler@uni-paderborn.de`

Abstract

This report documents the talks and discussions of Dagstuhl Seminar 14051 "Algorithms for Wireless Communication". The presented talks represent a wide spectrum of work on wireless networks. The topic of wireless communication continues to grow in many domains, new applications and deployments of wireless networks in a variety of contexts are being reported. A key focus of the talks and discussions presented here is to discuss models for wireless networks as well as algorithmic results and real world deployments.

Seminar January 26–31, 2014 – http://www.dagstuhl.de/14051

1998 ACM Subject Classification C.2.1 Network Architecture and Design: Wireless Communication, C.2.2 Network Protocols, C.4 Performance of Systems: Modelling Techniques, E.1 Data Structures: Graphs and Networks, F.2 Analysis of Algorithms and Problem Complexity, G.2.2 Graph Theory: Graph Algorithms

Keywords and phrases wireless, algorithms, model, complexity

Digital Object Identifier 10.4230/DagRep.4.1.152

1 Executive Summary

Guy Even
Magnús M. Halldórsson
Yvonne Anne Pignolet
Christian Scheideler

The last decades have seen an ever growing interest in wireless communication networks and their applications. Wireless networks pose many algorithmic challenges for various reasons: Realistic wireless signal propagation and interference models are very complex and therefore hard to use in rigorous algorithmic research, and this is further complicated by emerging technologies such as MIMO (multiple-input and multiple-output). Also, reasonable models for the dynamics and mobility in these networks can be quite complex and are not yet well-understood. Furthermore, standard complexity measures such as time and space are not sufficient any more as energy consumption is also a critical aspect that cannot be neglected. Many protocols for wireless networks have already been proposed by the research community, but most of them have only been studied in simulations or analyzed using rather simple models. So there is doubt whether any of these protocols would actually work in practice.

Algorithms for Wireless Communication, *Dagstuhl Reports*, Vol. 4, Issue 1, pp. 152–169
Editors: Guy Even, Magnús M. Halldórsson, Yvonne Anne Pignolet, and Christian Scheideler
DAGSTUHL REPORTS Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The purpose of this Dagstuhl seminar was to bring together computer scientists of different backgrounds to review and discuss models and algorithmic approaches in order to obtain a better understanding of the capabilities and limitations of modern wireless networks and to come up with more realistic models and approaches for future research on wireless networks that may then be investigated in joint research projects. The mix of the participating people resulted in fruitful discussions and interesting information exchange. The structure of the seminar took advantage of these different backgrounds by focusing on themed talks and open discussions.

The program included an eclectic mix of algorithmic and systems perspectives, modeling issues and emerging networking techniques, and explorations of the limits and possibilities of fundamental problems.

Discussions of models ranged from simple graph-based communication and interference models, to stochastic models, adversarial interruptions and jamming, dynamic networks and uncertainty formulations, and variations and extensions of signal-strength models.

Presentations from the systems perspective included managing environmental factors affecting measurements, robust predictions of channel capacities, efficiency of backpressure routing, issues in emerging heterogeneous radio environmental contexts, and robots controlled via wireless communication.

New dimensions at different networking layers included MIMO, network coding, interference cancellation, directional antennas and cognitive radio networks.

Finally, new results were presented on various related classic problems including broadcast, local broadcast, game theory, coding, routing, positioning, and connectivity.

2 Table of Contents

3 Overview of Talks

3.1 Jamming-Resistant Learning in Wireless Networks

Johannes Dams (RWTH Aachen, DE)

Joint work of Johannes Dams, Martin Hoefer and Thomas Kesselheim

We consider capacity maximization in wireless networks under adversarial interference conditions. There are n links, each consisting of a sender and a receiver, which repeatedly try to perform a successful transmission. In each time step, the success of attempted transmissions depends on interference conditions, which are captured by an interference model (e.g. the SINR model). Additionally, an adversarial jammer can render a $(1-\delta)$-fraction of time steps unsuccessful. Our main result is an algorithm based on no-regret learning converging to an $O(1/\delta)$-approximation. It provides even a constant-factor approximation when the jammer exactly blocks a $(1-\delta)$-fraction of time steps. In addition, we consider a stochastic jammer, for which we obtain a constant-factor approximation after a polynomial number of time steps. Using this learning approach and the general proof technique, we can even extend the results to more general settings, in which links arrive and depart dynamically, and where each sender tries to reach multiple receivers. Though these results cannot directly be applied to a setting with multiple channels and stochastic channel availabilities, we can achieve a constant-factor approximation using other learning approaches.

3.2 Braess Paradox in Wireless Networks: The Danger of Improved Technology

Michael Dinitz (Johns Hopkins University – Baltimore, US)

Joint work of Michael Dinitz and Merav Parter

When comparing new wireless technologies, it is common to consider the effect that they have on the capacity of the network (defined as the maximum number of simultaneously satisfiable links). For example, it has been shown that giving receivers the ability to do interference cancellation, or allowing transmitters to use power control, never decreases the capacity and can in certain cases increase it. But there is no reason to expect the optimal capacity to be realized in practice, particularly since maximizing the capacity is known to be NP-hard. In reality, we would expect links to behave as self- interested agents, and thus when introducing a new technology it makes more sense to compare the values reached at game-theoretic equilibria than the optimum values.

In this paper we initiate this line of work by comparing various notions of equilibria (particularly Nash equilibria and no-regret behavior) when using a supposedly "better" technology. We show a version of Braess Paradox for all of them: in certain networks, upgrading technology can actually make the equilibria worse, despite an increase in the capacity. We construct instances where this decrease is a constant factor for power control, interference cancellation, and improvements in the SINR threshold β, and is $O(\log n)$ when power control is combined with interference cancellation. However, we show that these examples are basically tight: the decrease is at most $O(1)$ for power control, interference cancellation, and improved β, and is at most $O(\log n)$ when power control is combined with interference cancellation.

3.3 Frequency Hopping against a Powerful Adversary

Yuval Emek (Technion, IL)

Joint work of Yuval Emek and Roger Wattenhofer

Frequency hopping is a central method in wireless communication, offering improved resistance to adversarial interference and interception attempts and easy non-coordinated control in dynamic environments. In this talk, we introduce a new model that supports a rigorous study of frequency hopping in adversarial settings. We then propose new frequency hopping protocols that allow a sender- receiver pair to exploit essentially the full communication capacity despite a powerful adversary that can scan and jam a significant amount of the ongoing transmissions.

3.4 Deterministic Rateless Codes for BSC

Guy Even (Tel Aviv University, IL)

Joint work of Guy Even, Benny Applebaum and Liron David

A rateless code encodes a finite length information word into an infinitely long codeword. The length of the prefix of the noisy codeword required by the decoder depends on signal-to-noise ratio of the communication channel. A rateless code achieves capacity for a family of channels if, for every channel in the family, reliable communication is obtained with a rate that is arbitrarily close to the channel's capacity. The encoder is universal because the same encoding is used for all channels in the family.

We construct the first *deterministic* rateless code for the binary symmetric channel. Our code can be encoded and decoded in $O(\log\log k)$ time per bit and in poly-logarithmic parallel time. Furthermore, the error probability of our code is almost exponentially small $\exp(\Omega(k/poly \log\log(k)))$. Previous rateless codes are probabilistic (i.e., based on code ensembles), require polynomial time per bit for decoding, and have inferior error probabilities.

3.5 Robot swarms as mobile sensor networks

Sándor Fekete (TU Braunschweig, DE)

Joint work of Becker, Aaron; Demaine, Erik; Fekete, Sándor; Habibi, Golnaz; Kroeller, Alexander; Lee, Soung Kyou; McLurkin, James; Schmidt, Christiane

We present two studies of algorithmic methods for swarms of mobile devices as sensor networks. In the first, we show a video describing exploration and mapping of an unknown environment by a swarm of robots with limited capabilities. In the second, we discuss positive and negative results for controlling a massive swarm of particles by uniform external forces in the presence of obstacles.

3.6 Arbitrary Transmission Power in the SINR Model

Fabian Fuchs (KIT – Karlsruhe Institute of Technology, DE)

Both in the light of energy conservation and the expansion of existing networks, wireless networks face the challenge of nodes with heterogeneous transmission power. However, for more realistic models of wireless communication only few algorithmic results are known. In this talk we propose to consider the transmission power as input. Thus we consider nodes with arbitrary power assignment in the so-called physical or SINR model. Our first result is a bound on the probabilistic interference from all simultaneously transmitting nodes on receivers. This result implies that current local broadcasting algorithms can be generalized to the case of non-uniform transmission power with minor changes. Also, we introduce a new network parameter ℓ that measures the length of the longest unidirectional path in the network. After showing that a dependence on ℓ is inevitable for distributed node coloring, we present a coloring algorithm its time complexity.

3.7 Modeling and Analysis of Wireless Networks using Stochastic Geometry

Martin Haenggi (University of Notre Dame, US)

Analyses of specific configurations of wireless networks are tedious and yield results without any generality. In contrast, wireless network ensembles described by spatial stochastic models are fairly tractable and give insight for a general class of networks. Stochastic geometry, in particular point process theory, provides these both these models and the mathematical tools for their analysis. The point process represents the locations of the nodes in a wireless networks. A prominent example is the Poisson point process, for which many interesting closed-form results have been derived. Originally used for sensor and ad hoc networks, it is now increasingly popular also as a model for cellular networks. While simple graphs do not constitute useful models for wireless networks, random geometric graphs based on point processes and the SINR link model are powerful and versatile performance indicators. Many metrics of interest can be extracted from such models, but it needs to be understood that many protocol decisions at the lower layers must be made before a meaningful graph can be defined. In other words, graphs should always be interpreted as the result of certain protocol choices, rather than as the starting point of an analysis or algorithm development.

3.8 Network Coding for Multi-Hop Wireless Broadcast

Bernhard Haeupler (Microsoft Research – Mountain View, US)

We introduce a simple distributed implementation of random linear network coding (RLNC) [8] and gives several scenarios arising in wireless network broadcast settings in which RLNC achieves large gains over traditional routing strategies.

After explaining an novel, simpler technique for analyzing RLNC [5] we survey several recent results obtaining the first throughput optimal algorithms in different wireless models using RLNC. In particular simple proof (sketches) for a throughput optimal broadcast in the radio network model [3, 4] and the dynamic network model [6, 2, 7] are presented.

Lastly we give some recent results on the network coding gap for broadcast in the radio network model [1].

References

1 N. Alon, M. Ghaffari, B. Haeupler, and M. Khabbazian. Broadcast throughput in radio networks: Routing vs. network coding. *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1831–1843, 2014.

2 C. Dutta, G. Pandurangan, Z. Sun, and E. Viola. On the Complexity of Information Spreading in Dynamic Networks. *ACM/SIAM Symposium on Discrete Algorithms (SODA)*, 2013.

3 M. Ghaffari and B. Haeupler. Fast structuring of radio networks large for multi-message communications. *International Symposium on Distributed Computing (DISC)*, 8205:492–506, 2013.

4 M. Ghaffari, B. Haeupler, and M. Khabbazian. Randomized broadcast in radio networks with collision detection. *ACM Symposium on Principles of Distributed Computing (PODC)*, pages 325–334, 2013.

5 B. Haeupler. Analyzing Network Coding Gossip Made Easy. *ACM Symposium on Theory of Computing (STOC)*, pages 293–302, 2011.

6 B. Haeupler and D. Karger. Faster Information Dissemination in Dynamic Networks via Network Coding. *ACM Symposium on Principles of Distributed Computing (PODC)*, pages 381–390, 2011.

7 B. Haeupler and F. Kuhn. Lower bounds on information dissemination in dynamic networks. *International Symposium on Distributed Computing (DISC)*, pages 166–180, 2012.

8 T. Ho, M. Médard, R. Koetter, D. R. Karger, M. Effros, J. Shi, and B. Leong. A random linear network coding approach to multicast. *IEEE Transactions on Information Theory (TransInf)*, 52(10):4413–4430, 2006.

3.9 Extending SINR to Realistic Environments

Magnús M. Halldórsson (Reykjavik University, IS)

Joint work of Marijke Bodlaender and Magnús M. Halldórsson

Signal-strength models of wireless communications capture the gradual fading of signals and the additivity of interference. As such, they are closer to reality than other models. However, nearly all theoretic work in the SINR model depends on the assumption of smooth geometric decay, one that is true in free space but is far off in actual environments. The challenge is to model realistic environments, including walls, obstacles, reflections and anisotropic antennas, without making the models algorithmically impractical or analytically intractable.

We present a simple solution that allows the modeling of arbitrary static situations by moving from geometry to arbitrary decay spaces. The complexity of a setting is captured by a “metricity” parameter ζ that indicates how far the decay space is from satisfying the triangular inequality. All results that hold in the SINR model in general metrics carry over to decay spaces, with the resulting time complexity and approximation depending on ζ in the same way that the original results depends on the path loss term α. For distributed algorithms, that to date have appeared to necessarily depend on the planarity, we indicate how they can be adapted to arbitrary decay spaces at a cost in time complexity that depends on a fading parameter of the decay space. In particular, for decay spaces that are doubling, the parameter is constant-bounded.

3.10 Online Independet Set with Stochastic Adversaries

Martin Hoefer (Universität des Saarlandes, DE)

Joint work of Goebel, Oliver; Hoefer, Martin; Kesselheim, Thomas; Schleiden, Thomas; Vöcking, Berthold

We investigate online algorithms for maximum (weight) independent set on graph classes with bounded inductive independence number ρ like interval and disk graphs with applications to, e.g., task scheduling, spectrum allocation and admission control. In the online setting, nodes of an unknown graph arrive one by one over time. An online algorithm has to decide whether an arriving node should be included into the independent set.

Traditional (worst-case) competitive analysis yields only devastating results. Hence, we conduct a stochastic analysis of the problem and introduce a generic sampling approach that allows to devise online algorithms for a variety of input models. It bridges between models of quite different nature – it covers the secretary model, in which an adversarial graph is presented in random order, and the prophet-inequality model, in which a randomly generated graph is presented in adversarial order.

Our first result is an online algorithm for maximum independent set with a competitive ratio of $O(\rho^2)$ in all considered models. It can be extended to maximum-weight independent set by losing only a factor of $O(\log n)$, with n denoting the (expected) number of nodes. This upper bound is complemented by a lower bound of $\Omega(\log n / \log^2 \log n)$ showing that our sampling approach achieves nearly the optimal competitive ratio in all considered models. In addition, we present various extensions, e.g., towards admission control in wireless networks under SINR constraints.

3.11 Multichannel Information Dissemination

Stephan Holzer (MIT – Cambridge, US)

Joint work of Stephan Holzer, Thomas Locher, Yvonne Anne Pignolet, Roger Wattenhofer

We study the information exchange problem on a set of multiple access channels: k arbitrary nodes have information they want to distribute to the entire network via a shared medium partitioned into channels. We present algorithms and lower bounds on the time and channel complexity for disseminating these k information items in a single-hop network of n nodes. More precisely, we devise a deterministic algorithm running in asymptotically optimal time $O(k)$ using $O(n^{\log(k)/k})$ channels if $k \leq \frac{1}{6}\log n$ and $O(\log^{1+\rho}(n/k))$ channels otherwise, where $\rho > 0$ is an arbitrarily small constant. In addition, we show that $\Omega(n^{\Omega(1/k)} + \log_k n)$ channels are necessary to achieve this time complexity.

3.12 Nearly Optimal Asynchronous Blind Rendezvous Algorithm for Cognitive Radio Networks

Qiang-Sheng Hua (Tsinghua University – Beijing, CN)

Rendezvous is a fundamental process in Cognitive Radio Networks, through which a user establishes a link to communicate with a neighbor on a common channel. Most previous solutions use either a central controller or a Common Control Channel (CCC) to simplify the problem, which are inflexible and vulnerable to faults and attacks. Some blind rendezvous algorithms have been proposed that rely on no centralization. Channel Hopping (CH) is a representative technique used in blind rendezvous, with which each user hops among the available channels according to a pre-defined sequence. However, no existing algorithms can work efficiently for both symmetric (both parties have the same set of channels) and asymmetric users. In this paper, we introduce a new notion called Disjoint Relaxed Difference Set (DRDS) and present a linear time constant approximation algorithm for its construction. Then based on the DRDS, we propose a distributed asynchronous algorithm that can achieve and guarantee fast rendezvous for both symmetric and asymmetric users. We also derive a lower bound for any algorithm using the CH technique. This lower bound shows that our proposed DRDS based distributed rendezvous algorithm is nearly optimal. Extensive simulation results corroborate our theoretical analysis.

3.13 Distributed Protocols for SINR

Tomasz Jurdziński (University of Wroclaw, PL)

In the advent of large-scale multi-hop wireless technologies, it is of utmost importance to devise efficient distributed protocols to provide basic communication. Unlike in the graph-based radio network model, algorithmic issues for multi-hop communication in the SINR model are not well understood. Especially, very few deterministic solutions are known.

The presentation addresses this issue for ad hoc SINR networks in the Euclidean space. Especially, the goal is to compare efficiency of determinstic and randomized solutions as well as to address the issue of the impact of knowledge of positions on complexity of basic communication problems.

3.14 Three wireless systems problems

Holger Karl (Universität Paderborn, DE)

The talk briefly highlights three problems in wireless communication. First, the streaming of video to mobile devices exploiting predictions of cellular capacity and user behavior. Second, the optimization of mobile backhaul networks to support wireless techniques like cooperative multipoint. Third, an optimization problem that combines cooperative and frequency diversity techniques, using OFDMA in a relay-extended cellular context.

3.15 Algorithmic Problems that Arise from Shifting to Directional Antennas

Matthew Katz (Ben Gurion University – Beer Sheva, IL)

Directional antennas have several noticeable advantages over omni-directional antennas. For example, when using a directional antenna, less power is required to transmit to a given receiver and less interference is caused by this transmission. Despite these advantages, the vast majority of the papers dealing with algorithmic problems motivated by wireless networks, consider omni-directional antennas, whose coverage area is often modeled by a disk. In this talk we discuss several basic algorithmic problems that arise when switching to directional antennas.

Let P be a set of points in the plane representing transceivers, and assume that each transceiver is equipped with a directional antenna. The coverage area of a directional antenna at point p of angle alpha, is a circular sector of angle alpha centered at p, where the orientation and range of the antenna can be adjusted. For a given assignment of orientations and ranges, the induced symmetric communication graph (SCG) of P is the undirected graph, in which there is an edge between two vertices (i.e., points) u and v if and only if v lies in u's sector and vice versa. The induced asymmetric communication graph (DCG) of P is the directed graph, in which there is a directed edge from u to v if and only if v lies in u's sector.

We consider several problems arising in wireless networks with directional antennas, under both the symmetric and asymmetric models, including orientation and range assignment, power assignment, and interference reduction. In the orientation and power assignment problem, one needs to assign an orientation and range to each of the antennas, so that the resulting communication graph is connected (in the symmetric model) or strongly connected (in the asymmetric model). The goal is to do so while keeping the ranges short.

The interference of a network is defined as the maximum in-degree of a node in DCG, i.e., the maximum number of transceivers covering a receiver. We address the following question under both models: What is the minimum interference I, such that for any set P of points in the plane, representing transceivers equipped with a directional antenna of angle alpha, one can assign orientations and ranges to the points in P, so that the induced communication graph G is either connected or strongly connected and the interference of G does not exceed I. We show that in the symmetric model, the advantage of directional antennas with respect to omni- directional antennas is less obvious.

3.16 Techniques for SINR Approximation Algorithms

Thomas Kesselheim (Cornell University, US)

We consider the capacity-maximization problem in the SINR model: Given pairs of senders and receivers, select a maximum subset of these such that all transmissions can be carried our simultaneously. For the problem variant with uniform powers, there is a constant-factor approximation due to Halldorsson and Mitra [SODA 2011]. We discuss a simplified analysis in spirit of [Kesselheim, ESA 2012]. The algorithm is similar in structure to the one for the problem variant where powers have to be chosen [Kesselheim, SODA 2011]. This allows us to derive an abstraction based on edge-weighted conflict graphs, which is useful for advanced problems.

3.17 Backpressure Routing in Wireless Networks: From Theory to Practice

Bhaskar Krishnamachari (University of Southern California, US)

Joint work of Bhaskar Krishnamachari, Scott Moeller, Majed Alresaini, Mike Neely, Andrea Gasparri, Shangxing Wang, Longbo Huang, Avinash Ridharan, and Omprakash Gnawali

Since the work by Tassiulas and Ephremides in 1992, Backpressure scheduling has been a focus of intense research by network control theorists. I describe our work on translating this theory into practice in the form of Backpressure Collection Protocol (BCP), the first implementation of dynamic bacpressure routing. I also discuss our work on backpressure with adaptive redundancy (BWAR), which utilizes multi-copy routing to improve the delay of backpressure routing for intermittently connected mobile networks. Finally, I present ongoing work on using backpressure scheduling to control the motion of message ferrying robots.

The work described was funded in part by the U.S. National Science Foundation via CNS-1049541.

3.18 Probabilistic Analysis of Power Assignments

Bodo Manthey (University of Twente, NL)

Joint work of Bodo Manthey and Maurits de Graaf

We consider the problem of assigning transmission powers to the devices of a wireless network such that the resulting communication graph is connected and the total transmit power is minimized. Our goal is a probabilistic analysis of this power assignment (PA) problem.

We prove complete convergence of PA for arbitrary combinations of the dimension d and the distance-power gradient p. In particular, we prove complete convergence for the case $p \geq d$. As far as we are aware, complete convergence for $p > d$ has not been proved yet for any Euclidean functional.

Furthermore, we prove that the expected approximation ratio of a simple spanning tree heuristic is strictly less than its worst-case ratio of 2.

3.19 Continuous Local Strategies for Robotic Formation Problems

Friedhelm Meyer auf der Heide (Universität Paderborn, DE)

Joint work of Bastian Degener, Barbara Kempkes and Peter Kling

I study a scenario in which n mobile robots with a limited viewing range are distributed in the plane and have to solve a formation problem. The formation problems considered are the GATHERING problem and the CHAIN- FORMATION problem. In the GATHERING problem the robots have to gather in one (not predefined) point, while in the CHAIN-FORMATION problem they have to form a connected chain of minimal length between two stationary base stations. Each robot has to base its decisions where to move only on the current relative positions of neighboring robots within its viewing range. Variants of these problems (especially for the GATHERING problem) have been studied extensively in different discrete time models. In this talk, I focus on a continuous time model where robots continuously sense the positions of other robots within their viewing range and continuously adapt their speed and direction according to some simple, local rules. Hereby, I assume that the robots have to obey a speed limit of one. I present strategies for both problems with $O(n)$ runtime bounds. Moreover, I show that their runtimes are at most by a factor $O(\log(OPT))$ and $O(\log(n))$, resp., away from the optimal time OPT necessary by a global algorithm, for GATHERING and CHAIN-FORMATION, resp.

3.20 Thoughts on Models for Wireless Networks

Calvin Newport (Georgetown University – Washington, US)

In my talk, I presented three thoughts about wireless algorithm research. The first thought was that it is useful to make a distinction between basic science research (exploring the capabilities and limits of wireless communication) and applied research (developing algorithms for use in practice). My second thought was that for applied research, it is often useful to introduce uncertainty to the relevant model. My third thought was that for applied research, we should also consider studying problems at the network layer, focusing less on contention (which is hard to formally model) and more on other aspects of wireless that show up at higher layers (uncertainty in delay and receivers, lack of network knowledge, etc.).

3.21 The topology of wireless communication and applications

Merav Parter (Weizmann Institute – Rehovot, IL)

Joint work of Chen Avin, Asaf Cohen, Yoram Haddad, Erez Kantor, Zvi Lotker, Merav Parter and David Peleg

We study the topological properties of wireless communication maps and their usability in algorithmic design. We consider the SINR model, which compares the received power of a signal at a receiver against the sum of strengths of other interfering signals plus background noise. To describe the behavior of a multi-station network, we use the convenient representation of a *reception map*. In the SINR model, the resulting *SINR diagram* partitions the plane into reception zones, one per station, and the complementary region of the plane where no station can be heard. We consider the general case where transmission energies are arbitrary (or non-uniform). Under that setting, the reception zones are not necessarily convex or even connected. This poses the algorithmic challenge of designing efficient point location techniques as well as the theoretical challenge of understanding the geometry of SINR diagrams. We achieve several results in both directions. One of our key results concerns the behavior of a (d+1)-dimensional map. Specifically, although the d-dimensional map might be highly fractured, drawing the map in one dimension higher "heals" the zones, which become connected. In addition, we study the topology of reception regions when reception points are allowed to decode messages using interference cancellation. Our final note concerns the connection between SINR diagrams and Voronoi diagrams.

3.22 Enhancing Future Networks with Radio Environmental Context

Marina Petrova (RWTH Aachen, DE)

As the demand for ever high data rates is continuing to increase driven by the global growth in smart phones and tablets, it is of great importance to re-think the spectrum allocation strategies and and at the same time enable cognitive radio technology to optimally use the available spectrum resources. Better understanding the radio environment in time and space and being able to dynamically predict interference could potentially help in providing capacity whenever needed and improving the wireless connectivity and service to the user. A Radio Environment Map (REM) is an advanced knowledge base that stores live multi-domain information on the entities in the network as well as the environment historically. The main functionality of a REM is the construction of dynamic interference map for each frequency at each location of interest by collecting spectrum measurements. By using advanced data processing methods and with a help of geographical terrain models, propagation environment, and regulations it can also estimate the state of locations where there is no measurement data. In this talk we will talk about the use of such REMs for predicting network coverage holes and discuss how using spatial statistics methods such as Kriging can produce reliable predictions under certain data accuracy and measurement density requirements.

3.23 Environmental-Aware Protocols for Networked Embedded Systems

Kay Römer (TU Graz, AT)

Sensor networks prodive 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 provides 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. In particular, environmental temperature deeply affects the operation of sensor networks. 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 applications as every new deployment site requires costly customization.

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, we first present TempLab, a testbed where user-defined temperature conditions can be created. For this, 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. Using this testbed, we systematically study the impact of temperature on the signal-to-noise ratio of a wireless link and provide a model that accurately predicts SNR for given transmitter and receiver temperatures. We also study the impact of temperature on the frequency of processor clocks, finding that a temperature change of 30 degrees slows down the processor clock by more than 10 percent. Finally we investigate the impact of temperature on the RPL routing protocol and find that the topology of the routing tree is significantly affected, even leading to network partitions.

3.24 Flow/back-pressure/slide approaches for routing and scheduling: from wired to wireless networks

Adi Rosén (University Paris-Diderot, FR)

We shortly review (older) results about flow/back-pressure/slide techniques for routing and scheduling in stable and dynamic networks. In particular we review results that show that such protocols, albeit being online and distributed, maintain stability in any network as long as the traffic injected into the network allows stability (roughly speaking, as long as there exist paths for the injected packets which imply injection rate at most 1). We then discuss the challenges of adapting these protocols to wireless networks and some perhaps promising directions for such adaptation.

3.25 SINR with adversarial noise

Christian Scheideler (Universität Paderborn, DE)

Joint work of Adrian Ogierman, Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang

In my talk I consider the problem of how to efficiently share a wireless medium which is subject to harsh external interference or even jamming. While this problem has already been studied intensively for simplistic single- hop or unit disk graph models, almost no work is known for the SINR interference model.

The talk consists of two parts. First, I introduce a new adversarial SINR model which captures a wide range of interference phenomena. Concretely, I consider a powerful, adaptive adversary which can jam nodes at arbitrary times and which is only limited by some energy budget. In the second part I present a distributed MAC protocol which provably achieves a constant competitive throughput in this environment: I show that, with high probability, the protocol ensures that a constant fraction of the non-blocked time periods is used for successful transmissions. The result also highlights an inherent difference between the SINR model and unit disk graph models.

3.26 The Power of MIMO

Christian Schindelhauer (Universität Freiburg, DE)

Joint work of Christian Schindelhauer and Thomas Janson

We present a wireless network model for MIMO, i.e. coordinated senders and receivers. First, we show that this model reduces to the standard SINR- model in the case of random senders and uncoordinated receivers. Then, we concentrate on the case of senders and receivers on a line. In our free-space model we coordinate the sender phase and amplitude in such a way that on the left we produce a beam while on the other side the path loss coefficient can be arbitrarily small. For equidistant ad-hoc nodes, which only coordinate themselves by the receipt times of messages we present a broadcasting algorithm with time $\mathcal{O}(\log n)$ for n nodes. In this setting the energy and SINR ratio only allows communication to the next node. So, cooperated to SINR we improve the time from $n-1$ to $\mathcal{O}(\log n)$.

3.27 Deterministic Blind Rendezvous in Cognitive Radio Networks

Ravi Sundaram (Northeastern University – Boston, US)

Joint work of Chen, Sixia; Russell, Alexander; Samanta, Abhishek; Sundaram, Ravi

Blind rendezvous is a fundamental problem in cognitive radio networks. The problem involves a collection of agents (radios) that wish to discover each other in the blind setting where there is no shared infrastructure and they initially have no knowledge of each other. Time is divided into discrete slots; spectrum is divided into discrete channels, $\{1, 2, ..., n\}$. Each agent may access a single channel in a single time slot and we say that two agents rendezvous when they access the same channel in the same time slot. The model is asymmetric: each agent A_i may only use a particular subset S_i of the channels and different agents may have access to different subsets of channels. The goal is to design deterministic channel hopping schedules for each agent so as to guarantee rendezvous between any pair of agents with overlapping channel sets. Two independent sets of authors, Shin et al. and Lin et al., gave the first constructions guaranteeing asynchronous blind rendezvous in $O(n^2)$ and $O(n^3)$ time, respectively. We present a substantially improved construction. Our results are the first that achieve nontrivial dependence on the size of the set of available channels. This allows us, for example, to save roughly a quadratic factor over the best previous results in the important case when channel subsets have constant size. We also achieve the best possible bound of O(1) time for the symmetric situation; previous works could do no better than O(n). Using the probabilistic method and Ramsey theory we provide evidence in support of our suspicion that our construction is asymptotically optimal for small size channel subsets.

3.28 Algorithmic Aspects of Geometric Routing on Mobile Ad Hoc Network

Takeshi Tokuyama (Tohoku University, JP)

In this talk, we discuss algorithms for geometric routing on a mobile ad hoc network. Although geometric routing on a static ad hoc network is well studied, the maintenance of the underlying network such as Delaunay graph causes difficulty if nodes move freely.

We present a simple and efficient strategy of geometric routing on moving nodes without much effort to update the underlying network. We also show the gap of theory and practice on mobile ad hoc network, and propose a research direction of geometric routing so that it can be used in practice.

Participants

- Eyjólfur Ingi Asgeirsson
Reykjavik University, IS
- Marijke Bodlaender
Reykjavik University, IS
- Johannes Dams
RWTH Aachen, DE
- Michael Dinitz
Johns Hopkins University – Baltimore, US
- Yuval Emek
Technion, IL
- Guy Even
Tel Aviv University, IL
- Sándor Fekete
TU Braunschweig, DE
- Fabian Fuchs
KIT – Karlsruhe Institute of Technology, DE
- Jie Gao
SUNY – Stony Brook, US
- Seth Gilbert
National Univ. of Singapore, SG
- Martin Haenggi
University of Notre Dame, US
- Bernhard Haeupler
Microsoft Corp. – Mountain View, US
- Magnus M. Halldorsson
Reykjavik University, IS
- Martin Hoefer
Universität des Saarlandes, DE
- Stephan Holzer
MIT – Cambridge, US
- Qiang-Sheng Hua
Tsinghua Univ. – Beijing, CN
- Thomas Janson
Universität Freiburg, DE
- Tomasz Jurdzinski
University of Wroclaw, PL
- Holger Karl
Universität Paderborn, DE
- Matthew J. Katz
Ben Gurion University – Beer Sheva, IL
- Thomas Kesselheim
Cornell University, US
- Christian Konrad
Reykjavik University, IS
- Bhaskar Krishnamachari
Univ. of Southern California, US
- Fabian Daniel Kuhn
Universität Freiburg, DE
- Bodo Manthey
University of Twente, NL
- Friedhelm Meyer auf der Heide
Universität Paderborn, DE
- Calvin Newport
Georgetown University – Washington, US
- Merav Parter
Weizmann Institute – Rehovot, IL
- Marina Petrova
RWTH Aachen, DE
- Yvonne-Anne Pignolet
ABB – Baden-Dättwil, CH
- Dror Rawitz
Bar-Ilan University, IL
- Kay Römer
TU Graz, AT
- Adi Rosén
University Paris-Diderot, FR
- Alexander Russell
University of Connecticut – Storrs, US
- Christian Scheideler
Universität Paderborn, DE
- Christian Schindelhauer
Universität Freiburg, DE
- Ravi Sundaram
Northeastern Univ. – Boston, US
- Takeshi Tokuyama
Tohoku University, JP
- Roger Wattenhofer
ETH Zürich, CH