



^b
UNIVERSITÄT
BERN

Proceedings of the WISH seminar on Wireless Integration of Sensor networks in Hybrid architectures

Editors: D.C. Dimitrova, G. Wagenknecht, M. Brogle

Technischer Bericht IAM-12-001 vom 10. April 2012

Institut für Informatik und angewandte Mathematik, www.iam.unibe.ch



Proceedings of the WISH seminar on Wireless Integration of Sensor networks in Hybrid architectures

Technischer Bericht IAM-12-001 vom 10. April 2012

CR Categories and Subject Descriptors:

C.2.1 [Computer-Communication Networks]: Network Architecture and Design;

C.2.4 [Computer-Communication Networks]: Distributed Systems

General Terms:

Design and Algorithms, Deployment, Measurement, Performance

Additional Key Words:

sensor networks

Institut für Informatik und angewandte Mathematik, Universität Bern

Abstract

The WISH (Wireless Integration of Sensor Networks in Hybrid Architectures) seminar was an opportunity for researchers from both academia and industry to share their work and views on the development in Wireless Sensor Networks (WSNs) and specifically their incorporation in hybrid networks. The seminar focused on research efforts at the territory of Switzerland. By bringing together academic researchers from different but related areas in contact with industrial peers, the seminar gave an opportunity to share expertise on ongoing research and stimulated the identification of future trends and collaboration possibilities.

The seminar program included ten technical talks, which provide a multidisciplinary view on WSNs and, in particular, on their deployment and practical value for variety of real-world applications. A major discussion topic was the deployment of sensor networks in real environments, e.g., for the purpose of precipitation monitoring, and the many challenges arising from that. Another strong focus was on the benefits of using sensor networks for enabling sustainable living. A third group of talks addressed the more technical aspects of developing a functional wireless sensor network, i.e., design, architectures and algorithms.

The technical talks provided a valuable feedback on what is being researched by academia, what is of interest for industry and what is the current state of match between the two areas. These proceedings contain the abstracts of the technical talks along with contact information of the authors.

Contents

1	Twilight: Twitter in Disaster Mode	1
2	Swiss Experiment. Merging high- and low-tech measurement data into the researcher's workflow	4
3	Sensor Networks for Diffusion Fields: Detection of Sources in Space and Time	7
4	Hobnet project	10
5	MOTEL A Mobile Robotic-Assisted Wireless Sensor Networks Testbed	13
6	Joint identity-message coding	16
7	Improving energy efficiency through opportunistic sensing	19
8	Online Multi-Model Approximation of Time Series	22
9	Benefits of wireless mesh networks for environmental sensing within the research project MontanAqua	25
10	Real-time spatiotemporal combination of radar and raingauge precipitation measurements in Switzerland	28

1 Twimight: Twitter in Disaster Mode

Paolo Carta

Communication Systems, ETHZ
E-mail: pcarta@ee.ethz.ch

Paolo Carta is a Research Engineer at the ETH Zurich. He is leading the Twimight internal project and the co-author of two publications. He holds an M.Sc with honours in Telecommunication Engineering from the University of Cagliari, Italy after completing his exchange at ETH Zurich.

Twimight:

Twitter in Disaster Mode

Theus Hossmann, Paolo Carta, Dominik Schatzmann, Franck Legendre
Communication Systems Group
ETH Zurich, Switzerland
firstname.lastname@tik.ee.ethz.ch

Per Gunningberg, Christian Rohner
Communication Research Group
Uppsala University, Sweden
firstname.lastname@it.uu.se

1. INTRODUCTION

In recent disaster events (e.g., earthquakes, floods), online social networks (OSNs) such as Facebook and Twitter have proven to be the communication tools of choice for affected people. However, these platforms were not explicitly designed for emergency situations in the first place. In fact, their value in such environments could be greatly improved by extending them with features specially designed for hostile circumstances.

One drawback of today's OSNs in disaster situations is their dependence on fixed network infrastructure. If the wired and/or cellular infrastructure is wiped out by natural forces (or even just congested in the aftermath of a disaster) the client-server communication of current OSNs is not feasible. Thus, in the time until emergency response forces are able to repair infrastructure or deploy temporary communication solutions (e.g., wireless mesh networks, satellite phones) people are left without any means to communicate.

To overcome this limitation, the use of delay tolerant opportunistic networks [1] has been proposed. Our goal is to augment OSN applications – which the users already have installed on their mobile phones and use in their every day life – with an additional feature, the *disaster mode*. In particular, we implemented the disaster mode in *Twimight*¹, our open source Twitter client. We choose Twitter since it has proven to be a highly useful communication platform in past emergency situations [2, 3]. Upon losing Internet connectivity, the user can simply enable the disaster mode to start *opportunistic communications* [4, 5].

Such mobility-assisted opportunistic communication has the advantage of being ready instantly after losing connectivity. No deployment of temporary infrastructure is required, since people have everything that is needed on their phones. However, it has the disadvantage of being a purely best-effort technology and comes with potential delays in tweet delivery – depending on the density of people and on their mobility. Hence, the disaster mode is mainly suited to provide a communication means to people in a *first phase* (during the hours or few days after a disaster happens), *before* emergency response gets active and reaches the disaster site. In a *second*




Figure 1: Two phases of disaster response.

and phase, there may be better suited technologies (e.g., satellite communication) to coordinate operations. We illustrate these two phases and the respective communication needs in Figure 1.

As a final note, we want to highlight that one of the main design goals of Twimight is *simplicity*. Hence, we do not aim at providing a highly specialized and sophisticated solution. Instead, we want to let people use the tools they know from their everyday life, modifying them as little as possible.

2. TWIMIGHT ARCHITECTURE

In this section we provide a brief overview of the architecture of Twimight with the main focus on the disaster mode and the opportunistic spreading of tweets². We have implemented *Twimight* as a disaster ready Twitter client for Android.

Twimight in Normal mode: *Twimight* supports most basic Twitter functionalities. In normal operation, *Twimight* queries Twitter for new tweets every 5 Minutes. The tweets obtained from the queries are cached locally in database.

Twimight in Disaster mode: *Twimight* has a checkbox in the settings menu to enable the disaster mode as shown in Figure 2b. The normal Twitter client functionalities stay the same but now rely on opportunistic communications. Upon enabling the disaster mode, the tweets are stored locally in a separate table for opportunistic spreading and later publication to the Twitter servers whenever connectivity is detected. Disaster tweets are highlighted in red on the user interface (Figure 2c) to mark them as important. The goal is that all tweets sent during a disaster can be received and displayed to as many people as possible.


In disaster mode, the device enables Bluetooth and scans period-

²For a more detailed discussion of Twimight functionality see also [6].


¹Download: <http://code.google.com/p/twimight>

ically for reachable Bluetooth devices. The scanning interval has been set to 2 minutes $\pm U[0, 20]$ seconds. Once two phones have discovered one another, they connect to each other and exchange the new disaster tweets, thereby spreading them epidemically.

From Disaster to Normal mode: As soon as connectivity to the Twitter servers is re-established, the disaster tweets sent by the user are published to the Twitter server. To ensure the reliable dissemination of tweets to the greater number, epidemic spreading (taking every opportunity to replicate a tweet) is mandatory at first. However, with the epidemic dissemination of tweets, not only do we potentially use a lot of network resources but victims of a disaster can also get overwhelmed with the number of received tweets that are displayed in their timeline. There are hence clear needs to scale such dissemination both at the network and human levels. In future work, we consider to limit the spreading of tweets to a given geographical range e.g., 5 km around the source. We also intend to leverage opportunistic routing protocols that avoid flooding all the network to reach the destination [7].



(a) Normal timeline and (b) Enabling disaster mode. tweet context menu.



(c) Timeline with high-lighted disaster tweet. (d) Sending tweets in disaster mode.

Figure 2: Twimight screenshots.

3. SPREADING OF SENSOR DATA OVER TWITTER

In this section we present a supporting application for the opportunistic Twitter: distributing sensor data with tweets, thus providing additional information about the user's environment. The

application runs autonomously without requiring user involvement and does not modify the opportunistic Twitter behaviour.

We extended Twimight to automatically tweet *sensor data* when in disaster mode. Sensor data can easily be formatted into tweets, do not require user interaction, and provide additional information about a user's environment.

Inferences about human activity, location and social events has been shown feasible on mobile phones by mobile and participatory sensing applications using microphone, accelerometer, gyroscope, camera, network interfaces, GPS, and other sensors built into or attached to modern mobile phones [8]. In our prototype we provide a stream of aggregated sensor readings from accelerometer and microphone with a moderate update interval of a few minutes. The sensor readings are aggregated and translated into activity states {motionless, stationary, going, running} and environment information {silence, voices, car, noisy, fire} respectively. The purpose to generate a periodic stream of sensor readings is to provide a sign of life. As a proof of concept, the states are decided upon statistical properties of the (low pass filtered) data.

The sensing is implemented as a Twimight plugin in the form of an Android Service that gets started when the disaster mode is enabled. Synchronization with Twimight is done using system-wide broadcast announcements. If location information is available on the phone (e.g., by the means of GPS readings), then the sensing tweets are geo-tagged as supported by the Twitter message format.

4. ABOUT AUTHORS

The Twimight project was initiated by Theus Hossmann (Postdoc at Cambridge) and Franck Legendre (Senior research at ETH Zurich) in collaboration with Christian Rohner and Per Gunningberg from Uppsala University. Paolo Carta is a Research Engineer at the ETH Zurich. He is leading the Twimight internal project and the co-author of two publications. He holds an M.Sc with honours in Telecommunication Engineering from the the University of Cagliari, Italy after completing his exchange at ETH Zurich.

5. REFERENCES

- [1] W. Zhao, M. Ammar, and E. Zegura, "A message ferrying approach for data delivery in sparse mobile ad hoc networks," in *ACM MobiHoc*, 2004.
- [2] A. Chowdhury, "Global pulse," <http://blog.twitter.com/2011/06/global-pulse.html>, 2011.
- [3] A. Bruns, "Tracking crises on twitter: Analysing #qldfloods and #eqnz." Presented at the Emergency Media and Public Affairs Conference, Canberra, 2011.
- [4] E. Nordström, P. Gunningberg, and C. Rohner, "A search-based network architecture for content-oriented and opportunistic communication," *Uppsala University Technical Report 2009-003*, 2009, <http://code.google.com/p/haggle>.
- [5] PodNet Project. [Online]. Available: <http://www.podnet.ee.ethz.ch/>
- [6] T. Hossmann, F. Legendre, P. Carta, P. Gunningberg, and C. Rohner, "Twitter in disaster mode: Opportunistic communication and distribution of sensor data in emergencies," in *ExtremeCom*, 2011.
- [7] T. Hossmann, T. Spyropoulos, and F. Legendre, "Know thy neighbor: Towards optimal mapping of contacts to social graphs for DTN routing," in *IEEE Infocom*, 2010.
- [8] H. Lu, J. Yang, Z. Liu, N. D. Lane, T. Choudhury, and A. T. Campbell, "The jigsaw continuous sensing engine for mobile phone applications," in *ACM SenSys*, 2010.

2 Swiss Experiment. Merging high- and low-tech measurement data into the researcher's workflow

Nicholas Dawes

Institut für Schnee- und Lawinenforschung
E-mail: dawes@slf.ch

Nicholas Dawe is currently leading research at the Swiss Institute for Snow and Avalanche Research (SLF) where he is manager of the Swiss Experiment Project (www.swiss-experiment.ch). As such he works on topics covering development of wireless infrastructures and in-line data processing as well as integration of the end system. Previously he has been an acoustic analyst and signal processing engineer at the Defence Science and Technology Laboratory, UK, and an avalanche scientist on at SLF, Switzerland.

Swiss Experiment

Merging high- and low-tech measurement data into the researcher's workflow

N. Dawes¹, S. Sarni², A. Salehi², H. Jeung², M. Bavay¹, A. Clifton¹, I. Iosifescu³, S. Michel², A. Wombacher², T. Egger⁴, K. Aberer², M. Lehning¹.

1. WSL / SLF, Flüelastrasse 11, CH-7260 Davos Dorf
2. LSIR, EPFL, BC 109 (Bâtiment BC), Station 14, CH-1015 Lausanne
3. IKG, Wolfgang-Pauli-Strasse 15, CH-8093 Zürich
4. Egger Consulting, Vienna, Austria

In order to study the full range of processes occurring in our environment, we not only need to study all of the disciplines and look at effects on multiple timescales, but we also need to study the interdisciplinary processes and the inter-scalar relationships. The development of cheaper sensors and self-organising networks has made it possible to study many effects over wide areas with high spatial and temporal resolution. These measurements improve our understanding of the processes occurring, but they still cannot cover all scales and all disciplines. Experiments are carried out for a specific purpose, in a specific place, but it is only when we make these measurements available together with measurements from all of the disciplines at multiple fieldsites, both nationally and internationally, that we can really start to perform the types of interdisciplinary and inter-scalar science that we truly need.

The Swiss Experiment Platform (SwissEx) project is an initiative of the Competence Centre Environment and Sustainability (CCES). The primary motivation behind SwissEx is to provide the technology for a common, cross-disciplinary platform where data and metadata from all disciplines at a variety of temporal and spatial resolutions are accessible.

Within SwissEx, new technologies are under development that will enable widespread use of dense spatial measurements where cost and/or resources have previously been a barrier. These range from deployment methods for single sensors such as the Distributed Temperature Sensor (DTS), to wireless sensor networks and new, low powered sensors which can be used within these WSNs. To complement the new sensor technologies, we have provided a distributed data infrastructure for a generic environmental science usecase. This infrastructure aims to use the best technologies to store and share both metadata and data along with automated data quality recognition. Interfaces are under development to combine all of this information into a package that will allow scientists to use data in an informed manner. These interfaces will allow data to be accessed from anywhere, at any time and will allow the integration of data access into 3rd party interfaces.

A generic infrastructure can provide great advantages, such as:

- greater cross-disciplinary visibility of available data (data are often used for multiple disciplines)
- cross-institution availability of master datasets (read only for non-authors)
- reduction of data ownership issues (local storage and access management)
- access to multiple datasets using a single tool (the location of data storage becomes irrelevant to the user)

The development of such an infrastructure, however, brings unique challenges not faced by other infrastructures built around specific sensors or scientific disciplines:

- visualisations cannot be optimised according to specific parameters or their context
- distributed systems which are optimised for scientists to manage their own data inherently create multiple data managers with varying data management experience
- any centralised data portal can create a bottleneck which would quickly be crippled by multiple users downloading large datasets.
- if a standard ontology is not imposed (which is difficult across diverse communities), databases rapidly lose their interoperability.

Multi-disciplinary data systems provide excellent tools for locating data, but most eventually provide a series of local files for further processing, providing marginal advantages for frequent users. We have integrated a web-service infrastructure and the plugin tools to be able to query data over this web-service infrastructure from commonly used processing tools, which are already integrated into the scientists workflow. With this system, we have enabled the researcher to import data from both his/her own experiments together with other applicable data which may be stored anywhere in the world. They can then immediately start working with the data, just by entering a few lines of code.

In a new project starting in September 2012, we will take the existing SwissEx components, align them with other projects worldwide and take the infrastructure forwards to become an environmental knowledgebase. This knowledgebase will build upon some of the ideas already available within SwissEx, e.g. the addition of contextual information to the dataset by linking publications and methods, such that the likelihood of data misinterpretation (and hence conflicting publications) may be reduced, as well as building up a library of measurement methods and information on environmental sensing. Together with a variety of users, we will also build up a library of generic processing tools with connected interfaces to the data and information on their application. These tools will utilise the latest knowledge in data control, spatio-temporal statistics, trend analysis and downscaling. The data discovery aspect of SwissEx will be improved by greater integration of our GIS platform, including joint environmental data/map data queries, better handling of spatially distributed data (3 and 4 dimensional), and improved tools for data browsing (quick look tools).

We present this generic architecture together with examples of WSNs integrated with existing sparse spatial measurement infrastructures.

3 Sensor Networks for Diffusion Fields: Detection of Sources in Space and Time

Ivan Dokmanic

Audiovisual communications Lab, EPFL

E-mail: ivan.dokmanic@epfl.ch

Ivan Dokmanic received the Dipl. Ing. degree in electrical engineering from the University of Zagreb, Croatia, in 2007. From 2007 to 2010 he was a teaching assistant at the Faculty of Electrical Engineering and Computing, University of Zagreb. From 2006 to 2007 he was also working as a codec developer for MainConcept AG, and as a digital audio effects developer from 2009 to 2010 at Little Endian Ltd. Currently he is working towards a PhD degree in the Audiovisual Communications Laboratory (LCAV) at EPFL. His research interests include inverse problems in acoustics, digital audio processing, signal processing of physical fields and environmental signal processing. He received the Best student paper award at ICASSP 2011.

Sensor Networks for Diffusion Fields: Detection of Sources in Space and Time

Ivan Dokmanić, EPFL-LCAV

I. INTRODUCTION

We present some recent research going on in the Audio-visual Communications Laboratory (LCAV) on sensing and reconstruction of the diffusion fields. In this initial investigation we concentrate on the diffusion and diffusion-like processes since these models describe many environmentally relevant phenomena. A common purpose of the presented algorithms is learning about the sources of a physical field from measurements collected by a sensor network.

Recovering the sources of a physical field is an *inverse problem*. Assuming that we know the generative model, source reconstruction enables us to compute the field in all space and time. Traditional sampling and reconstruction of signals is also an inverse problem. Thus if the signals are physical fields, we have two sets of constraints: (1) the samples and (2) the partial differential equation (PDE) describing the field. Blending the two approaches, we engineer algorithms for spatio-temporal field reconstruction. Ultimately, one would like to produce a single *black box* formalism which solves the problem invariantly of the underlying physical phenomenon.

In this work, we are primarily concerned with the sampling and the reconstruction of environmentally relevant physical fields, but the proposed algorithms are not specific to these scenarios.

We are interested in problems such as the detection of the pollution source, the reconstruction of the time-varying plume source, or the estimation of the temperature distribution. These phenomena are modeled by the diffusion equation,

$$\Delta u(\mathbf{x}, t) - \frac{1}{D} \frac{\partial u(\mathbf{x}, t)}{\partial t} = f(\mathbf{x}, t), \quad (1)$$

with $f(\mathbf{x}, t)$ being the source distribution.

In all of above applications the field sources are spatially sparse. Taking this to the limit, we use the concept of a point source, which we represent by the Dirac delta function. Concretely,

$$f(\mathbf{x}, t) = \sum_{k=1}^K s_k(t) \delta(\mathbf{x} - \mathbf{x}_k), \quad (2)$$

with s_k modeling the temporal variation of k th source. If we also assume that $s_k(t) = c_k \delta(t - t_k)$, then we model sudden events in time, such as explosions.

The authors are with the School of Computer and Communication Sciences, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland (e-mails: ivan.dokmanic, martin.vetterli@epfl.ch). Martin Vetterli is also with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720, USA.

This work was supported by an ERC Advanced Grant – Support for Frontier Research – SPARSAM Nr: 247006.




Fig. 1. Illustration of the problem setup. We show how to reconstruct the diffusion field driven by K instantaneous sources from spatio-temporal samples (how to go from the left hand side to the right hand side of the figure).

Since in practice a pure diffusion is rare, we coin the term *diffusion-like processes*. By diffusion-like, we think of the processes that either exhibit a strong diffusion component or that have a diffusion-like decay. For instance, we might talk about the diffusion-advection fields, where in addition to pure diffusion we also have the transport of the particles by the medium. Such equations model the dispersion of a plume generated by a smokestack—a point source.

We will analyze two different theoretical problems. The first problem is sensing and reconstruction of the diffusion fields [1]. We consider the fields driven by point sources in space and time, and we are interested in retrieving the source parameters—release times, magnitudes and locations. This scenario is illustrated in Fig. 1. For simplicity we show the theory in the 1-D case noting the possibility of a multi-dimensional extension. The second problem concerns the monitoring of the emissions from multiple smokestacks using a sensor network. Key observation is that we know the locations of smokestacks, but want to indirectly measure their temporal variation. We approach it by assuming that the emission waveforms conform to certain low-dimensional models.

II. PROBLEM 1: DIFFUSION SAMPLING

First, we analyze the inverse problem of reconstructing a diffusion field such as temperature from samples collected by a sensor network. Inverse problems are often ill-conditioned, meaning that small errors in the measured data can lead to large errors in the solution. The reconstruction of the diffusion fields is known to be particularly ill-conditioned, thus requiring strong assumptions on the source distribution. In our case we assume that individual sources are concentrated to spatio-temporal points.




Fig. 2. A sketch of the sensing scenario. The M smokestacks are at known locations $\{\xi_j\}$. The sensors are represented as blue circles and are located at $\{\zeta_i\}$. We estimate the flow of substance that is released by each smokestack $s_j(t)$.

We observe that the terms in the eigenfunction expansion of the field decay rapidly. This is a consequence of a specific eigenvalue distribution for the diffusion equation. Thus, we approximate the field by a truncated series and show that the approximation error decays rapidly with time. On the other hand, the useful *information content* in the field also decays with time, suggesting the need for a proper choice of the sampling strategy. We propose two algorithms for sampling and reconstruction of the field. The first one reconstructs the distribution of point sources appearing at known times using the finite rate of innovation (FRI) framework. The second algorithm addresses a more difficult problem of estimating the unknown times at which the point sources appear, in addition to their locations and magnitudes. It relies on the assumption that the sources appear at distinct times. We verify that the algorithms are capable of reconstructing the field accurately through a set of numerical experiments. Specifically, we show that the second algorithm successfully recovers an arbitrary number of sources with unknown distinct release times.

The proposed diffusion field sampling involves three tunable parameters: spatial sampling frequency, temporal sampling frequency, and the cutoff index N . If we sample too late after the source had appeared, the signal-to-noise ratio (SNR) is too low to make any reasonable inference from these samples. On the other hand, if we sample too close to the source (in time and space), the bandwidth is large. Therefore, for a fixed N , the truncated approximation may be inaccurate, leading to the failure of the reconstruction algorithm. This shows that we should choose N according to the desired spatial and temporal sampling frequency.

III. PROBLEM 2: EMISSION MONITORING

The second problem we study is the spatio-temporal sampling of physical fields representing the dispersion of a polluting substance in the atmosphere. We consider the following setup: N sensors are deployed at the ground level and measure the concentration of a specific substance emitted by M smokestacks. This scenario is illustrated in Fig. 2.

The transport of the substance in the atmosphere is mainly the result of three processes: advection by the wind, diffusion (partially from turbulent eddy motion), and gravitational settling. We assume that we can measure the substance concentration in different spatio-temporal points using an opportunely designed sensor network. This scenario may be of significant

importance in the design of *citizen sensing* projects, such as OpenSense and SafeCast, and in enforcing environmental regulations.

Smokestack emissions vary over time, and it is this variation that we aim to infer. To recover the emission rates of the smokestacks with a limited number of spatio-temporal samples, we assume that the emission waveforms live on one of the following two low-dimensional structures,

- i) **Model 1:** The waveforms $s_j(t)$ belong to the subspace spanned by K known functions $\phi_{k,j}(t)$:

$$s_j(t) = \sum_{k=1}^K \alpha_{k,j} \phi_{k,j}(t), \quad (3)$$

where the $\alpha_{k,j}$ are the unknowns of the j th source.

- ii) **Model 2:** The waveforms $s_j(t)$ belong to a class of signals with the finite rate of innovation (FRI). Namely, each smokestack produces only K innovations over the considered period of time. This abstraction efficiently models many realistic signals including piecewise constant or piecewise polynomial emissions.

The assumed source models are important since they are sufficiently flexible to deal with many types of sources, and they provide an elegant way to solve the problem of estimating their appearance times. In other words, they effectively regularize the otherwise ill-conditioned inverse problem. We propose efficient algorithms and sufficient conditions for the recovery of the emission rates.

Finally, the techniques presented here can be applied to other sensing scenarios. Namely, if we consider any physical field modeled by a linear partial differential equation, we can recover the emission rates of its sources from the measurements collected by a sensor network, using the proposed algorithms.

REFERENCES

- [1] I. Dokmanic, J. Ranieri, A. Chebira, and M. Vetterli, "Sensor Networks for Diffusion Fields: Detection of Sources in Space and Time," in *Proc. Allerton Conf. on Comm., Control and Comp.*, 2011, pp. 1552–1558.
- [2] J. Ranieri, I. Dokmanic, A. Chebira, and M. Vetterli, "Sampling and Reconstruction of Time-Varying Atmospheric Emissions," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Proc.*, Oct. 2012, pp. 1–4.

4 Hobnet project

Orestis Evangelatos

TCS Sensor Lab, University of Geneva
E-mail: Orestis.Evangelatos@unige.ch

Orestis Evangelatos is currently a PhD student in the University of Geneva, Switzerland. He is a member of the Theoretical Computer Science and Sensor Networks Lab (TCS&SN). He received his diploma (equiv. MSc) in the Computer Engineering and Informatics Department in the University of Patras (Greece). His main interests are wireless sensor networks, energy efficient algorithms, routing, obstacle avoidance, dynamic networks, smart buildings and radiation aware. He is a teaching assistant in the Computer Science department of UniGE and he has contributed to the conception and realization of both European and national research projects.

HOBNET presentation

Orestis Napoleon Evangelatos

Department of Informatics

University of Geneva

Geneva, Switzerland

orestis.evangelatos@unige.ch

Abstract — This document is an abstract of the presentation of the HOBNET project that will be given in the WISH seminar on March 15th in Bern.

Key words: *HOBNET, wireless sensor networks, energy efficiency, building management, IPv6, 6LoWPan, CoAP*

I. INTRODUCTION

The HOBNET project stands for Holistic Platform Design for Smart Buildings of the Future InterNET. It is a STREP Research Project under the FIRE program. Its main aim is “to ease and maximize the use of FIRE platforms for Future Internet Applications on automation and energy efficiency for smart/green buildings”. Several partners from all around Europe contribute in the realization of the project. The companies of Ericsson Serbia (Serbia), Mandat International (Switzerland), Sensinode (Finland) and the universities of Patras (CTI-Greece), Edinburgh (Scotland), College Dublin (Ireland) and university of Geneva are combining their research skills in order to achieve the greater goal of the project. The project started on June 2010, whereas now we are in the middle of its duration as it will be completed on May 2013.

II. MAIN AIM

The main aim of the project includes research of algorithmic, networking and application development aspects of Future Internet systems of tiny embedded devices. Including but not limited to, the main aims concern:

- a) *An all IPv6/6LoWPAN network infrastructure* of buildings and how IPv6 can integrate heterogeneous technology (sensors, actuators, mobile devices etc)
- b) *6lowApp and its standarization* towards a new embedded application protocol for building automation
- c) *Novel algorithmic models and scalable solutions* for energy efficiency and radiation-awareness, data dissemination, localization and mobility
- d) *Rapid development and integration of building management applications*
- e) *Support for the deployment and monitoring* of resulting applications on FIRE testbeds.

III. METHODOLOGICAL APPROACH

We take a holistic approach addressing critical aspects at different layers in an integrated way, including the following hierarchy:

- At the low level, network protocols and architectures, mainly based on IPv6, are studied, with emphasis on heterogeneity and interoperability.
- At the second layer, we provide algorithmic models and solutions for smart buildings, with special care for scalability.
- At the third level, an interface for the rapid development and evaluation of building management applications is provided.
- Finally, proposed research solutions and key innovations are organically evaluated in the context of the platform integration.

IV. OBJECTIVES

1) Objective 1

Identification of the needs, current limitations and novel research challenges towards developing smart/green building environments. Proposal of guidelines for large scale deployments of sensors and actuators networks in buildings.

2) Objective 2

Design and evaluation of scalable all IPv6/6LoWPAN network architectures to support future internet devices and applications, particularly for the smart/green building management domain. The creation of a simulation framework for networks of heterogeneous IP networks.

3) Objective 3

Come up with a coherent set of novel models and implemented, tested and validated best algorithmic solutions and high level technical recommendations particularly for smart building scenarios.

4) Objective 4

Development of a service oriented architecture easing the mutual integration of several control and monitoring systems. Provide an interface layer between the building management system and FIRE experimentation platforms to be used for the rapid development and the evaluation of building management applications. The overall goal is to ease and maximize the utility of FIRE for the research and deployment of real building management systems (BMS).

5) Objective 5

Implementation, integration and evaluation of a platform prototype including proposed research solutions and key innovations. A broader research goal is to test the potential of IPv6 to deploy and integrate heterogeneous sensors, including

non IP based sensors together with 6LoWPAN sensors, actuators and mobile devices.

6) Objective 6

Contribute to 6LoWPAN and its standardization towards a new embedded application protocol for building automation. Development of a multipurpose building automation demonstration integrated into the project website.

V. INNOVATION

Rather than an application-agnostic infrastructure, the project addresses the specific R&D area of intelligent Building Management Systems (BMS). It's interoperable as it's interconnecting a variety of wireless devices (sensors, actuators, RFID tags, mobile phones etc.) as well as various hardware and software types for each device. The holistic approach to future internet systems of tiny embedded devices and the overall platform integration, in contrast to most existing approaches which focus on particular layers is again an innovation of the project. There is a lot of contribution to the standardization of the 6LoWPAN, especially towards a new embedded application protocol and application commissioning. Deep examination of scalability and mobility across large remote buildings is concern, not just a smart home or a single building as usually in the state of the art.

VI. ARCHITECTURE

The overall architecture of the project is based on a two domain concept:

Building Domain: Includes building automation devices, local services and a building resource directory. Main part of this component is the Open Building Interface (OBI).

Application Domain: includes local and remote building management applications and an application resource directory where its core component is the Embedded Building Interface (EBI).

The interaction between these two domains is handled by the Building Web Service Proxy (BWSP), which provides data services for the Application Domain and handles the interaction with Building Domain Hobnet compliant resources. The resource directory concept is a key piece of the Hobnet's architecture, with all the resources available to other components in a domain being registered and searchable via a resource directory.

In the EBI layer are placed the Hobnet managed IPv6 Devices. This zone includes the IPv6 sensor-actuator networks and devices that form the building domain of the Hobnet architecture. The devices specified in this zone are the gateway devices, the 6LoWPAN networking devices and the ZigBee networking devices. Focusing a little bit in the 6LoWPAN devices we could define the following for the OSI layers:

- 1) *Physical layer:* IEEE 802.15.4 (868 MHz, 2.4GHz).
- 2) *MAC layer:* IEEE 802.15.4 standard
- 3) *Network layer:* 6LoWPAN communication achieved via routed data in 6LoWPAN and IETF ROLL RPL (routing protocol for low power and lossy networks)

4) *Transport layer:* UDP protocol, which is commonly used in 6LoWPAN networks

5) *Application layer:* *Constrained Application Protocol* (CoAP) protocol. The EBI interface will specify application level interactions around the CoAP

VII. TESTBEDS

All the research and the all the achievements of the project are going to be implemented and experimentally validated in 3 testbeds located in UniGe, UniPatras and in Mandat. Currently the two first are already set up, running and testing algorithms. Some of the scenarios are already implemented in our testbeds. The testbed of Geneva is a heterogeneous sensor network testbed comprised by iSense and TelosB motes. For the iSense motes the WISELIB is used while for the programming of TelosB motes is used TinyOS and Contiki. Through the Contiki OS we can already control the deployed sensors using IPv6. For the moment this is done in a local area network but in the future the motes will be controlled by the Hobnet's infrastructure over the internet. Hobnet makes use of the IETF standard Web Linking and CoRE Link Format for describing resources in resource registration and lookup operations. The resource registration interface is used to allow a CoAP or HTTP end-point to register its resources in a resource directory. The resource registration interface has been contributed to the IETF CoRE. The resource registration is performed simply as sending a POST request to the RD with an appropriate query string defining the name of the end-point and lifetime of registration. Other requests are PUT, GET and DELETE.

VIII. TARGET USERS AND BENEFITS

Future Internet research and competitiveness in Europe can benefit by HOBNET's study of critical open issues like the interoperability of different networking technology, the integration of IPv6 with sensor networks of various types, the algorithmic scalability and the development of new standards for the application layer. HOBNET's Service Oriented Architecture (SOA) for Building Management Systems (BMSs) can be exploited by the Building Sector European Industry as well as by Public Utility National Activities on Green Buildings. As one of the aims of the project is the movement towards green buildings, the next generations of ICT could benefit in order to support lower carbon emissions for better energy efficiency, lightning, and more efficient environmental simulations and monitoring. From the scope of the designing and development, researchers could be able to test their high level algorithms in hardware (not just simulation) at a large scale, in realistic scenarios, while engineers would benefit from the interaction with rigorous algorithmic methodologies. On the other side the algorithmic and distributed computing community will benefit from the definition of more realistic abstract models and well-motivated problems for sensor networking.

Geneva, 24.2.2012

5 MOTEL A Mobile Robotic-Assisted Wireless Sensor Networks Testbed

Kamini Garg

NetLab, SUPSI

E-mail: kamini.garg@supsi.ch

Kamini Garg is a Ph.D. candidate at NetLab, SUPSI. She holds a Master's degree in Information Technology from Indian Institute of Information Technology and Management, India. Some of her recent work has been published in international conferences and workshops such as EWSN, MobiOpp and Wintech. She is actively participating in the MOTEL project.

MOTEL — A Mobile Robotic-Assisted Wireless Sensor Networks Testbed

Alexander Förster*, Anna Förster†, Kamini Garg†, Daniele Puccinelli†, Silvia Giordano† and Luca M. Gambardella*

*Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA), Switzerland

†Networking Laboratory, ISIN-DTI, University of Applied Sciences of Southern Switzerland

Abstract—We present MOTEL: a robotic-assisted mobile wireless sensor network testbed. It is able to deploy and conduct mobile WSN experiments, where sensor nodes are piggybacked on mobile robots. The system consists of two main components: Multi-Robot Architecture for Coordinated Mobility (MuRobA) and Flexible WSN Runtime Management Software Architecture (FLEXOR). The first controls the mobility of the robots and the second controls the sensor nodes without the use of a backchannel. In this demonstration, we show the general architectures and usage of both components. Most importantly, we show the deployment and usage of MOTEL as whole, which can be conducted in any indoor environment in only few hours.

I. INTRODUCTION AND MOTIVATION

Wireless sensor network testbeds have developed to be the de facto standard for testing WSN applications and algorithms. However, most of these testbeds rely on a complex infrastructure (backchannel) to provide power to the nodes and to disseminate code and data. This infrastructure is not only costly, but makes the testbed rigid. The nodes cannot be moved between or during experiments and re-deployment of the complete testbed is time and effort-consuming.

On the other side, mobile WSN testbeds offer a new environment to the WSN developer. However, it exhibits two major challenges: First, the backchannel disappears and second, the mobility of the nodes needs to be implemented. The problem of the backchannel has been addressed many times in terms of remote re-programming or debugging of nodes, e.g. the tool Marionette [1] or the Contiki toolchain [2]. However, such tools are rather rigid and support only one operating system (TinyOS or Contiki), while we are looking for a general-purpose lightweight user-controlled tool, which enables the testbed user to implement its own debugging commands and to easily exchange software modules at runtime.

The second challenge, the mobility of sensor nodes, has been usually addressed by piggybacking sensor nodes on large indoor robots. This has been implemented in the Wisebed platform as RoombaNet [3] or in the integrated CONET testbed (*conet.us.es*). However, all these systems use autonomous robots, which rely on complex localization and navigation algorithms, which often prove to be unreliable and slow. On the other side, there exist also centralized solutions, called also *global vision*, for mobile robots, like the ones used in the RoboCup small size league [4]. Here, the robots are observed by overhead cameras and localization and navigation becomes fast and reliable. This suits the purposes of a WSN testbed

much better and is also less costly and more scalable. In the next paragraphs, we present our solutions to both problems: FLEXOR and MuRobA and will demonstrate their deployment and usage in MOTEL.

II. FLEXOR: FLEXIBLE RUNTIME MANAGEMENT SOFTWARE ARCHITECTURE FOR WSN

FLEXOR is a general-use software architecture for programming WSNs. It is platform-independent and has extensive graphical support for implementing, programming and managing WSNs. Its general architecture is detailed in [5] and depicted also in Figure 1. For MOTEL, the most important properties of FLEXOR are its general-use remote function call mechanism (implemented by the Callback Manager) and the possibility to exchange software components at run-time. For the latter, we define *Images*, which consists of several *specifications*, which in turn describe software components architectures. For example, one specification might consist of an application and routing modules and another of the same application, but different routing module. At run-time, these specifications can be exchanged with a single 1-byte command instead of complex re-programming and reboot of the complete system.

FLEXOR provides the following further functionalities and abilities to MOTEL:

- Remote function call of user-defined functions
- Parameter change of individual modules
- Runtime exchange of software modules
- Remote debugging, status inquiries and data logging

All together, these properties enable MOTEL to run sophisticated, structured experiments with WSNs in both mobile and static environments without the need of a backchannel. In order to enable the mobility of the nodes, we piggyback them on mobile robots and implement the robotic architecture MuRobA, described briefly in the next section.

III. MUROBA: MULTI-ROBOT ARCHITECTURE FOR COORDINATED MOBILITY

The general architecture of MuRobA is presented in Figure 2. It consists of one to several cameras, overlooking the robots; the component *camview*, which tracks the colorful dots on top of the robots to localize them; the *FleetManager*, which consolidates the information of all cameras and decides the movement commands for all robots; and finally one or several




Fig. 2. MuRobA architectural overview, with its main components the camera input, the fleet manager and the bluetooth-supported robot control. Note that the system can have several cameras as input and several bluetooth controllers as output.




Fig. 1. FLEXOR system overview with its main components. Most important for MOTEL is its remote function calls (CallbackManager) and the possibility to exchange software components at run-time (the specifications).

BluetoothTransmitter, which sends commands to the robots via bluetooth. Note that the system is very flexible and scalable, as it allows for several camera inputs and several bluetooth controllers. Note also the modularization of the system makes it hardware independent. It allows for the usage of any cameras with any lenses and any robotic platforms.

IV. OVERVIEW OF MOTEL

Figure 3 depicts the MOTEL deployment. It consists of the playground with the camera overlooking it, the robots with sensor nodes piggybacked on them, and the two control stations for MuRobA and FLEXOR. Additional sensor nodes can be freely placed wherever it fits the experiment and will be also controlled through FLEXOR.

V. NEXT STEPS

In the next future, we will implement and enable a remote testbed control unit, so that experiments can be planned and conducted also remotely on MOTEL. This is a major challenge especially because of the batteries needed by both robots and sensor nodes. We plan to supply the sensor nodes with power through the robots and to implement hardware and software for the robots, which will enable them to re-charge autonomously.




Fig. 3. General architecture of MOTEL with its main components for controlling the robots and the sensor nodes.

ACKNOWLEDGMENTS

The work presented in this paper has been funded partially by SUPSI under grant number 32RAX3MOTEL, "MOTEL: Mobile wireless sensor network testbed" and partially by the EU Cooperating Objects Network of Excellence (CONET).

REFERENCES

- [1] K. Whitehouse, G. Tolle, J. Taneja, C. Sharp, S. Kim, J. Jeong, J. Hui, P. Dutta, and D. Culler, "Marionette: using rpc for interactive development and debugging of wireless embedded networks," in *The Proceedings of the 5th International Conference on Information Processing in Sensor Networks*, 2006. *IPSN 2006*, 2006, pp. 416–423.
- [2] A. Dunkels, N. Finne, J. Eriksson, and T. Voigt, "Run-time dynamic linking for reprogramming wireless sensor networks," in *Proceedings of the 4th ACM International Conference on Embedded Networked Sensing Systems (SenSys)*, Boulder, CO, USA, 2006.
- [3] M. A. Hail, J. Pinkowski, T. Teubler, M. Danckwardt, D. Pfisterer, and H. Hellbrück, "Roombanet - testbed for mobile networks," in *Proceedings of the Workshops der wissenschaftlichen Konferenz Kommunikation in verteilten Systemen 2011 (WowKiVS 2011)*, T. Margaria, J. Padberg, and G. Taentzer, Eds., vol. 37. Electronic Communications of the EASST, 2011, accepted for publication.
- [4] A. Egorova, A. Glove, C. Göktekin, A. Liers, M. Luft, R. Rojas, M. Simon, O. Tenchio, and F. Wiesel, "Robocup 2004 symposium: Papers and team description papers," in *RoboCup 2004 Symposium: Papers and Team Description Papers*, 2004.
- [5] A. Förster, K. Garg, D. Puccinelli, and S. Giordano, "Flexor: User friendly wireless sensor network development and deployment," 2011, under submission.

6 Joint identity-message coding

Lorenzo Keller

Algorithmic Research and Networked Information, EPFL
E-mail: lorenzo.keller@epfl.ch

Lorenzo Keller is pursuing a Ph.D. in the School of Computer and Communication Sciences, EPFL, Switzerland. His research mainly focuses on wireless networks. In particular he is interested in studying how to integrate ideas coming from theory, such as network coding, in practical systems.

Joint identity-message coding

Lorenzo Keller
EPFL, Lausanne, Switzerland

Abstract—In typical sensor network applications it is necessary to know the identity of the sensor that generated each collected measurement. To do so traditionally the identity of the source is appended to each measurement as a separate field when the measurement is sent. In this presentation we show that by jointly coding messages and identities it is possible to create packets that can be efficiently compressed at nodes that route large amount of them and therefore reduce the energy they consume to transmit. By reducing the maximal energy consumption in the network this approach can significantly reduce costs when the network topology cannot be known at deployment time and therefore all nodes are provisioned to support the worst-case energy consumption.

I. INTRODUCTION

In many sensor network applications both measurements and identity of their source must be collected. A typical example of such type of application is when a sensor network is used to reconstruct a spatial field of a physical quantity such as temperature. Often the measurements can be represented with a small number of bits. For instance when measuring temperatures the changes from one measurement to the next can be represented with few bits if the sampling frequency is high enough.

Traditionally the identity of the source of a packet is attached as an header to the data and it is a negligible overhead. For instance in IP networks the source identity is represented with 4 bytes while a typical payload is around 1500 bytes. However when the data to be collected is very small and the number of sources is very large, as it is often the case in sensor networks, the size of this header becomes a significant portion of the transmitted traffic.

Reducing the overhead due to identity headers is particularly compelling on nodes that have to forward large amounts of traffic. In nodes with low traffic the measurement headers are a small overhead compared to link layer headers, on nodes with high traffic the overhead of link layer headers is amortized by sending multiple measurements in the same packet and therefore the source identity headers become a significant fraction of the bytes transmitted.

On high traffic nodes it is possible to reduce the overhead due to source identity headers by compressing many of them together. Even if measurement data is not compressible the identity headers are. For instance if a node is forwarding the measurement of all sensors of the network it can simply transmit the measurements sorted by source identifier: in this case the identity of the source of each measurement is implicitly represented by its position in the packet.

This presentation is based on the work published in [3]

In this presentation we present an approach that allows to perform such compression in a straightforward way. Each source node maps its measurement to a codeword. Multiple source codewords can be compressed by simply XORing them bit by bit. The resulting compressed codewords can be further combined using XOR at successive routers. The codewords are designed in such way that the sink is able to decompress them and obtain the measurements and corresponding source identities. The approach is based on the concept of subspace coding[4]. This presentation describes a simplified version of the codes, for a full analysis please refer to [3].

The proposed coding scheme reduces the energy consumption on the most busy nodes. This is very important if such nodes cannot be identified at deployment time and therefore all node must be provisioned with enough energy to support that function. The downside of this approach is that the codewords generated by the sources are larger than a packet encoded with a traditional header and payload scheme therefore they increase the energy consumption of the nodes that do not forward much traffic. To limit this effect we design codes that can be adapted to situations in which not all the source measurements need to be compressed in the same packet.

II. PROTOCOL DESCRIPTION

We consider a network composed of N sensor nodes $\sigma_1, \dots, \sigma_N$ and a sink. The nodes are organized in a tree rooted at the sink. The network operates in rounds. At each round each sensor s_i chooses a message from a set of possible message $\mathcal{M} = \{1, \dots, M\}$ and forms a tuple $t_i = \langle i, m_i \rangle$ where i is a unique identifier of the node and m_i is the chosen message. For simplicity in the following we will assume that $M = 2^\Delta$. Sources map their tuple to a codeword as explained in Section III. Leaves of the tree simply forward the codeword to their parent. Each interior node combines codewords generated locally and received from the children using bit-by-bit XOR. At most K codewords are combined together. To prevent combination of more than K sources an header counting the number of combined codewords is appended to each packet. Nodes then forward each of the resulting combined codewords to their parent. At the end of each round the sink has collected combined codewords from its children and can decode the tuples chosen by the sources.

III. CODING SCHEME

Each sensor σ_i maps its tuple $t_i = \langle i, m_i \rangle$ to a bit vector v_{i,m_i} of length l bits. The tuples are associated to vectors as follows: Let $\mathbf{H}$ be a $l \times N\Delta$ parity check matrix of a




Fig. 1. Comparison of the average codeword length for different protocols. The performance of coding is bounded by the two lines represented in the plot.

binary error correcting code of length $N\Delta$, redundancy l and minimum distance $\min\{2K\Delta + 1, N + 1\}$. Partition the columns of $\mathbf{H}$ in N sets $C_1, \dots, C_N$ of Δ columns. For each sensor σ_i map each of the 2^Δ subsets $V_{i,1}, \dots, V_{i,M}$ of C_i to one of the M messages. The codeword for message j is then $v_{i,j} = \bigoplus_{v \in V_{i,j}} v$.

To show that the sink can distinguish packet formed from different set of tuples we use a property of error correcting codes: given a linear code with minimum distance d_{min} , any $d_{min} - 1$ columns of the parity check matrix $\mathbf{H}$ are linearly independent [5], and therefore their sum is non-zero.

Let v_T a packet received by the sink which is the sum of the codewords corresponding to the tuples in set T . Let V_T the set of columns of $\mathbf{H}$ used to build the codewords used by the sensors to send the set of tuples T . Let T, T' any two distinct sets of tuples. In order for the sink to decode we need the following: $v_T \neq v_{T'} \Leftrightarrow \bigoplus_{v \in V_T} v \neq \bigoplus_{v \in V_{T'}} v \Leftrightarrow \bigoplus_{v \in V} v \neq 0$ where $V = (V_T \cup V_{T'}) - (V_T \cap V_{T'})$. Observe that since $T \neq T'$ then $V_T \neq V_{T'}$ and therefore $|V| > 0$ and since $|V_T \cup V_{T'}| < d_{min}$ then $|V| < d_{min}$. The sum of columns of $\mathbf{H}$ in V therefore cannot be zero. This proves that the sink to correctly distinguish between different sets of tuples. The sink can efficiently decode the received combinations using the approach described in [2].

IV. CODING SCHEME PERFORMANCE

In order to quantify the benefits of the proposed coding scheme we compute the average number of bits necessary to represent a set of measurements averaged over all possible sets of up to K measurements. We compare four coding schemes: aggregation, the traditional approach that separates identity from payload and simply concatenates multiple measurements, our coding scheme when using codewords that allows to combine up to K measurements, our coding scheme when using codewords that allow to combine all measurements and the optimal code that maps any of the possible sets of up to K measurements to a variable length bit-vector such that the average length is minimal. Notice that for our scheme




Fig. 2. Number of bytes transmitted by the busiest node including link layer headers. Since aggregated packets sometimes exceed the MTU they need to be fragmented, the part of bytes belonging to additional fragments is indicated as subsequent packets.

we propose a low complexity compression and decompression algorithm which is not available for the optimal scheme.

For lack of space we refer the reader to [3] for the details of the analysis. Figure 1 shows the performance of the different coding schemes. The proposed scheme significantly outperforms aggregation as soon as many measurements are combined and has a performance that is almost optimal when the payload is very small.

V. PROTOCOL PERFORMANCE

To assess the performance of the protocol in a realistic setup we implemented it as a TinyOS module and we tested it with TOSSIM. Our implementation extends the Collection Tree Protocol[1], the standard collection protocol of TinyOS, to use our coding scheme. Children are automatically discovered and a mechanism is in place to recover from tree changes. For more details about our implementation and the test setup refer to [3].

Figure 2 compares the number of bytes sent by the busiest node in the network when simply aggregating multiple identity-message pairs and when using our coding scheme (in this case using $K = N$). The network is a square grid of nodes, each node is well connected to four of its neighbors while it is only marginally connected to the other nodes, the sink is in one of the corners of the grid. We can see that our coding scheme dramatically reduces the amount of bytes sent by the busiest node. More measurements are reported in [3], in particular the average number of bytes transmitted increases by less than 50%.

REFERENCES

- [1] Omprakash Gnawali, Rodrigo Fonseca, Kyle Jamieson, David Moss, and Philip Levis. Collection tree protocol. In *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems, SenSys '09*, pages 1–14, New York, NY, USA, 2009. ACM.
- [2] M. Jafari Siavoshani, L. Keller, C. Fragouli, and K. Argyraki. Compressed network coding vectors. In *Information Theory, 2009. ISIT 2009. IEEE International Symposium on*, pages 109–113, 28 2009-july 3 2009.
- [3] L. Keller, M. Jafari Siavoshani, C. Fragouli, K. Argyraki, and S. Diggavi. Joint identity-message coding. *Selected Areas in Communications, IEEE Journal on*, 28(7):1083–1093, september 2010.
- [4] Ralf Koetter and Frank R. Kschischang. Coding for errors and erasures in random network coding. In *Information Theory, 2007. ISIT 2007. IEEE International Symposium on*, pages 791–795, june 2007.
- [5] F. J. MacWilliams and N. J. A. Sloane. *The Theory of Error-Correcting Codes*. North-Holland Mathematical Library, 1977.

7 Improving energy efficiency through opportunistic sensing

Wilhelm Kleiminger

Distributed Systems, ETHZ
E-mail: wilhelmk@ethz.ch

Wilhelm Kleiminger obtained his MEng (Computing) degree from Imperial College London in 2010. In his Master's thesis he worked on stream processing under the supervision of Dr Peter Pietzuch. Since 2010 he is a doctoral candidate at ETH Zurich, where he works on opportunistic sensing and smart energy. He is currently working on deriving occupancy information from opportunistic sensors such as smart phones and smart electricity meters in order to integrate this information into a smart heating system.

Improving energy efficiency through opportunistic sensing

Wilhelm Kleiminger
Institute for Pervasive Computing
ETH Zurich, Switzerland
Email: wilhelmk@ethz.ch

Abstract—We introduce the concept of an intelligent thermostat based on opportunistic sensing of household occupancy. By utilising Web-enabled sensors we aim to overcome the setup overhead of conventional systems and allow for dynamic integration of new sensor information. In the final system we aim to use prediction to control a thermostat with a probabilistic schedule in order to achieve the correct temperature upon the arrival of occupants.

I. INTRODUCTION

As smart phones become ever more widely used and the Internet of Things advances into people's homes, Information and Communication Technology (ICT) become a vital means to improve energy efficiency. ICT helps to increase efficiency by monitoring the usage of appliances and providing feedback to the users. By notifying users about the greatest energy guzzlers in the household, we can create an awareness in consumers that results in an overall reduction of energy consumption [7].

However, in order to achieve substantial savings, we have to go beyond visualisation and engagement strategies. Savings diminish if users are not continuously involved. User participation is difficult to sustain over longer periods of time if no appropriate incentives are given. The savings from programmable thermostats for example are often only realised if the occupants are generally inclined to save energy in the first place [9], [10]. In our work we focus on smart heating control algorithms which employ automation where possible and user interaction where necessary. To this end we propose to utilise ubiquitous sensors such as smart phones, smart electricity meters and other household sensors opportunistically to build occupancy schedules. Such a system may use, depending on the availability of the data, the current position of the occupant using GPS or connections to WiFi access points in order to determine where a person is and when she will most likely return home.

In the following sections we will describe in detail how our smart heating system can be used to provide significant energy savings in private households.

II. AN OPEN SENSING INFRASTRUCTURE

The number of sensors in households is constantly increasing. Modern smart phones, which are equipped with 3-dimensional accelerometers, differential GPS and near-field-communication (NFC), can be used to obtain information




Figure 1. A possible opportunistic sensing scenario with a smart phone, a wireless router, a smart power outlet, a smart meter and an infrared (PIR) sensor.

about the occupants' routines. By observing WiFi logins from a wireless router we can deduce whether the occupant has returned home and therefore build a schedule for the household occupancy. This way of data acquisition from existing sensors is usually referred to as *opportunistic sensing* [2]. By utilising already existing hardware to gather sensor data the initial cost per household is reduced.

Figure 1 shows a household instrumented with Web-enabled sensors. Traditionally, the use of Internet protocols has been regarded as too resource-intensive for embedded devices [12]. Sensor data was collected through proprietary and application-specific protocols. Recent work by Hui et al. showed that it is possible to implement IP on resource-constrained devices [5]. This has been leveraged to use conventional Web protocols to allow for using sensors and actuators for physical mashups [3].

Weiss et al. have demonstrated how Web technologies can be used to connect an electricity meter to a smart phone in order to make users aware of their the electricity consumption [13]. We plan to build upon this work and access additional opportunistic sensors to construct our smart heating system.

III. SMART HEATING CHALLENGES

Off-the-shelf programmable thermostats ask the user to enter their schedule in order to fall back to a setback temperature

when the home is unoccupied. For properties with multiple occupants, a common schedule must be found and entered. If the occupants change their schedules, this must reflect in re-programming of the thermostat. Moreover, the users must be aware of the time it takes for the home to heat up to the comfort temperature. If this ramp-up time is not taken into account when programming the thermostat, the home might not have sufficient time to heat up to a comfortable temperature.

Recent work has shown that the challenges of re-programming the thermostat, including the computation of the ramp-up time can be best solved by a smart heating system that senses and predicts occupancy from sensor values. Lu et al. showed how a thermostat can be augmented with reed switches and infrared sensors to sense the arrival times of occupants and program the thermostat with a predicted schedule [6]. Gupta et al. show how the ubiquity of GPS sensors in smart phones can be used in conjunction with the predicted travel time to adjust the temperature just in time for the arrival of occupants [4].

IV. AN OPPORTUNISTIC APPROACH

Our system improves on the existing work by opportunistically leveraging existing sensors such as smart electricity meters and smart phones. An intelligent thermostat that depends only on GPS traces from smart phones will always be limited by battery constraints. On the other hand, if we were to use only the smart electricity meter to sense occupancy, we might characterise the reading of a book on a bright winter day as absence from home and therefore lower the temperature. Likewise, the spikes in consumption generated by an electric boiler might inadvertently make the system believe that the occupants have returned home. A similar reasoning applies to inferring an unoccupied property from a lack of movements detected by an infrared sensor. We aim to reduce these false positives by utilising the different sensing modalities. If the system is not sure about the current state of the household as sensed by the smart electricity meter, it may ask for more information from the occupant's smart phone such as the last GPS coordinates. By utilising all currently available sensors on demand we try to build a more reliable smart heating system.

Furthermore, the user-input itself should serve as a means to improve as well as evaluate the efficiency of the system. As users provide feedback on the predicted schedule by manually overriding the temperature when they arrive in a cold house, we can establish a measure of comfort and evaluate how likely the occupants are to sacrifice some comfort in order to increase energy savings. This is crucial in households with a high variance of arrival times as this has a direct impact on the expected miss time of the prediction algorithm.

V. DATA COLLECTION

We are currently working on a deployment in 6 Swiss households during which we want to test opportunistic sensing for occupancy detection. Figure 1 shows the setup envisaged. During the experiment we will record electricity consumption data from a smart electricity meter and a number of

smart power outlets. In contrast to previous smart metering deployments, the data will be sampled at a rate of 1 Hz, allowing for a detailed analysis of the electricity consumption. In addition, we will collect data from infrared sensors and ground truth occupancy data through an interactive logbook on a tablet computer. The tablet computer will be installed near the entrance of the building in order for occupants to record their arrival and departure times.

Using the ground truth occupancy information, we aim to test an occupancy classification algorithm for smart meter data. We also want to show how the accuracy of the classification may be improved as additional information such as firings from infrared sensors or the disaggregated consumption figures from the smart power outlets are included.

REFERENCES

- [1] Validating the impact of programmable thermostats. RLW Analytics, Middletown, Connecticut, January 2007.
- [2] Andrew T. Campbell, Shane B. Eisenman, Nicholas D. Lane, Emiliano Miluzzo, Ronald A. Peterson, Hong Lu, Xiao Zheng, Mirco Musolesi, Krist Fodor, and Gahng-Seop Ahn. The Rise of People-Centric Sensing. *Internet Computing, IEEE*, 12(4):12–21, July/August 2008.
- [3] D. Guinard, V. Trifa, and E. Wilde. A Resource Oriented Architecture for the Web of Things. In *Internet of Things (IOT 2010)*, pages 1–8. IEEE, 2010.
- [4] Manu Gupta, Stephen S. Intille, and Kent Larson. Adding GPS-Control to Traditional Thermostats: An Exploration of Potential Energy Savings and Design Challenges. In *Proceedings of the 7th International Conference on Pervasive Computing (Pervasive 2009)*, pages 1–18. Springer, May 2009.
- [5] J.W. Hui and D.E. Culler. IP is dead, long live IP for wireless sensor networks. In *Proceedings of the 6th ACM conference on Embedded network sensor systems*, pages 15–28, New York, New York, USA, November 2008. ACM.
- [6] Jiakang Lu, Tamim Sookoor, Vijay Srinivasan, Ge Gao, Brian Holben, John Stankovic, Eric Field, and Kamin Whitehouse. The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems (SenSys 2010)*, pages 211–224. ACM, November 2010.
- [7] Friedemann Mattern, Thorsten Staake, and Markus Weiss. Ict for green – how computers can help us to conserve energy. In *Proceedings of the ACM e-Energy Conference*, April 2010.
- [8] Simon Mayer and Dominique Guinard. An Extensible Discovery Service for Smart Things. In *Proceedings of the 2nd International Workshop on the Web of Things (WoT 2011)*, pages 7–12, June 2011.
- [9] Monica J. Nevius and Scott Pigg. Programmable thermostats that go berserk: Taking a social perspective on space heating in wisconsin. ACEEE, August 2000.
- [10] Daniel Perry, Cecilia Aragon, Alan Meier, Therese Peffer, and Marco Pritoni. Making energy savings easier: Usability metrics for thermostats. *Journal of Usability Studies*, 6(4):226 – 244, August 2011.
- [11] Marla Sanchez, Richard Brown, Greg Homan, and Carrie Webber. Savings estimates for the United States Environmental Protection Agency's ENERGY STAR voluntary product labeling program. Environmental Energy Technologies Division, University of California, Berkeley, June 2008.
- [12] Jean-Philippe Vasseur and Adam Dunkels. *Interconnecting Smart Objects with IP - The Next Internet*. Morgan Kaufmann, 2010.
- [13] Markus Weiss, Tobias Graml, Thorsten Staake, Friedemann Mattern, and Elgar Fleisch. Handy feedback: Connecting smart meters with mobile phones. In *Proceedings of MUM 2009 (ACM International Conference on Mobile and Ubiquitous Multimedia)*, Cambridge, UK, November 2009.

8 Online Multi-Model Approximation of Time Series

Athanasios Papaioannou

Distributed Information Systems, EPFL
E-mail: thanasis.papaioannou@epfl.ch

Thanasis G. Papaioannou is a postdoctoral fellow at the Distributed Information Systems Laboratory of Ecole Polytechnique Fédérale de Lausanne (EPFL). He received his B.Sc. (1998) and M.Sc. (2000) in Networks and in Parallel/Distributed Systems from the Department of Computer Science, University of Crete, Greece, and his Ph.D. (2007) from the Department of Computer Science, Athens University of Economics and Business (AUEB). From spring 2007 to spring 2008, he was a Visiting Professor in the Department of Computer Science of Athens University of Economics and Business. From 1998, he has worked for 13 European projects, including MONTAGE, M3I, MMAPPS, GridEcon, Hydrosys, Wattalyst and OpenIoT, as well as Swiss projects SwissExperiment and OpenSense. He also has a significant publication record (over 30 papers) to high quality journals and conferences including Springer Electronic Commerce Research, Elsevier Computer Networks Journal, IEEE INFOCOM, IEEE ICDE, ACM SoCC, IEEE CCGRID etc. His research interests are in data stream processing, mechanism design for online environments, cloud resource management, privacy, trust and reputation, sensor networks, and QoS.

Online Multi-Model Approximation of Time Series

Thanasis G. Papaioannou, Mehdi Riahi, and Karl Aberer

School of Computer and Communication Sciences

Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland

Email: firstname.lastname@epfl.ch

Abstract—Real-time model-based data approximation and filtering is a common solution for reducing the storage (and communication) overhead. However, the selection of the most efficient model depends on the potentially dynamic characteristics of the data stream, namely rate, burstiness, deviation, etc. Here, we investigate the innovative concept of efficiently combining multiple approximation models in real-time for maximizing achievable data compression.

I. INTRODUCTION

Recent advances in sensor technology have enabled the availability of a multitude of (often privately-held) sensors. Embedded sensing functionality (e.g. noise, accelerometer, temperature, GPS, RFID etc.) is now included in mobile devices, such as phones, cars, buses, etc. Environmental and health-care applications based on community sensing in urban areas have been already envisioned, e.g. personalized carbon exposure and impact calculators, healthy lifestyle estimators, traffic monitoring, etc. The large amount of these (mobile) sensing devices and the huge volume of raw monitored data pose new challenges for the sustainable storage and efficient retrieval of sensor data streams.

Among different time series approximation techniques, piecewise linear approximation has been widely used [1], [2] for the online approximation of time series within a certain error norm (i.e. *lossy* compression). These models exploit the inherent correlations (e.g. with time or among data streams) in time series to split data in pieces and approximate each segment with a certain mathematical function derived by the model. However, the potential varying burstiness (and possibly rate) of the data streams along time and the variable standard error introduced by the sensor mobility often result in limited effectiveness of a single model for approximating data within the prescribed error bound during a certain period. The same argument is also valid for other time series that may exploit variable burstiness in different time periods, e.g. stock prices during volatile or non-volatile market periods. For example, as illustrated in Fig. 1, Swing linear model approximates better than the Piecewise Midrange Constant (MidRange) one the beginning of the stream, while the situation is reversed at the end of the stream period.

In this abstract (based on [3], which was presented at IEEE MDM'11), we propose the innovative concept of combining multiple statistical models for approximating time series. The intuition behind this approach is that different data periods of a sensor stream can be better approximated by different models, thus resulting, overall, in fewer and longer segments. This is




Fig. 1. Different data segments are better approximated by different models.

because of the greater flexibility offered by the alternative models. We propose our multi-model longest-fit algorithm and prove its correctness for approximating the data stream within the specified error bound. By an extensive series of experiments with both real and artificial data traces, we prove that our approach always produces fewer or equal segments than any of its constituent models employed individually.

II. THE APPROACH

Online piecewise approximation algorithms seek to find the parameters of a certain mathematical function, so as to fit the raw values of a segment of the data stream within a maximum error bound. When a raw data value cannot be approximated within the error bound by a specific instantiation of the “fitting” function of the model, then a new data segment is initiated; within the new data segment, a new instantiation of the fitting function has to be found and employed for data approximation. Each approximation model has to employ a fixed number of initial raw values in a data segment, in order to find an instantiation of its fitting function for this segment; e.g., 1 value for the cache filter, 2 values for the linear filter, 3 values for 2nd-degree polynomial regression, etc. We refer to this model requirement as *initialization*. The raw values of the segment necessarily fit into the model function during its initialization; the minimum initialization length that may be required by a model is 1 value. Note that, after the initialization, the model parameters may still change as new data arrives to update the instantiation of the model.

In our approach, a collection of models are jointly employed for approximating the data stream. Each data segment is approximated by the most effective model instantiation for that segment. The model effectiveness for a segment is determined based on the *segment length* (i.e. fitting period) in terms of raw data values that are approximated by the same instance of the

model and its achievable *compression ratio* for this segment.

More formally, we want to select the model m that better approximates each segment i of the data stream, so that the total number N of segments is minimized. There is no known optimal online algorithm for finding the most efficient combination of models. To this end, we propose a greedy algorithm for selecting for each segment the model that i) maximizes its length (i.e. approximates the largest number of raw values), and ii) it is the cheapest to be stored. The algorithm achieves this as follows: Consider a set M of models that jointly approximate a certain data segment of the data stream S . Each raw data item $\langle t, v \rangle$, i.e. with value v at time t , is examined by each of the initialized models for this data segment whether it falls (“hit”) or not (“miss”) within the error bound ϵ from the estimated data value by the model at time t , i.e. whether $|f_m(t) - v| < \epsilon$ or not for a model m with an instantiated function f_m . All uninitialized models succeed into approximating the raw data item $\langle t, v \rangle$ in this segment by default. We calculate the compression ratio (i.e. storage cost of uncompressed data over storage cost of model) r_m for each missing model m , and we then exclude m from the models that are further examined against the aforementioned *hitting condition* for this segment. We repeat examining the hitting condition for all subsequent raw data items of the stream, until all models “miss”. At this point, the model m^* with the highest compression ratio r_{m^*} is dumped into the database to approximate the data segment until the time t_{m^*} that it missed. Afterwards, the data stream is retracted to time t_{m^*} (i.e. the data stream is considered from time t_{m^*} onwards), the approximation formulas of all models are cleared and all models are considered for the approximation of the next data segment according to the aforementioned procedure. If at the time that all models miss, there are several missed models with the same highest compression ratio, then the one with minimum root mean squared error is selected to be dumped in the database.

The optimal offline combination of arbitrary models could be found by Dijkstra algorithm for finding the shortest path in communication networks with complexity $O(|S||M| + |S| \log |S|)$ for a data stream S and a set of models M .

III. EVALUATION

In our experiments, both real and synthetic data sets were employed (cf. [3] for the full set of results). As real data sets, we used measurements for various environmental parameters (air temperature, humidity, wind direction) collected from existing sensor deployments in the Swiss Alps by the Swiss Experiment project (www.swiss-experiment.ch) and ocean surface temperature sensor data from the TAO project (www.pmel.noaa.gov/tao/). Air temperature, ocean surface and humidity time series have smooth statistical behavior (due to their inherent physical laws), while the wind direction data set is highly *bursty* and quite unpredictable in nature. All data sets were approximated within 3 different maximum error bounds: 3.16%, 5%, and 10% of the data range. We implemented multiple models including Swing (SW), MidRange (MR),

Linear Filter (LF), Linear Regression (LR), Least Squares Line (LS), Constant Filter (CF), and Chebyshev Polynomial with different degrees (referred to as *Cheb* in figures with the degree specified in parenthesis).

We only present here the assessment of the compression effectiveness of our multi-model approximation as compared to the compression achieved by its constituent linear models when individually applied to the real data sets. Combining linear models is interesting, because they are computationally very efficient, very cheap to store, and they achieve comparable effectiveness to their more complex counterparts. As depicted in Figs. 2(a) and 2(b), our multi-model algorithm achieves better compression ratio, than any individual linear model, both for smooth (humidity) and bursty (wind direction) real time series respectively. However, the more bursty the data set, the lower the achievable compression ratio. Moreover, as experimentally proved, different linear models are selected by the multi-model approximation algorithm for approximating the different data segments of the stream. For example, the segments of the humidity stream are approximated 43% by SW, 30% by LS, 19% by MR, 6% by LR and 2% by CF.




Fig. 2. Compression ratios for a) Humidity and b) Wind Direction datasets.

IV. CONCLUSION

We investigated the innovative concept of combining multiple models that approximate time series within a maximum error bound for achieving higher compression effectiveness than the individual models themselves and experimentally found that 80% compression improvement can be achieved. Moreover, our greedy algorithm is very close ($\leq 6\%$) to the offline optimal achievable multi-model compression.

REFERENCES

- [1] I. Lazaridis and S. Mehrotra, “Capturing sensor-generated time series with quality guarantees,” in *Proc. of ICDE*, March 2003, pp. 429 – 440.
- [2] H. Elmeleegy, A. K. Elmagarmid, E. Cecchet, W. G. Aref, and W. Zwaenepoel, “Online piece-wise linear approximation of numerical streams with precision guarantees,” in *Proc. of VLDB*, Lyon, France, August 2009.
- [3] T. G. Papaioannou, M. Riahi, and K. Aberer, “Towards online multi-model approximation of time series,” in *Proc. of IEEE Mobile Data Management (MDM)*, Lulea, Sweden, June 2011.

9 Benefits of wireless mesh networks for environmental sensing within the research project MontanAqua

Bruno Schädler

Hydrology Group, University of Bern
E-mail: bruno.schaedler@giub.unibe.ch

Dr. Bruno Schdler completed a PhD thesis on the topic "The variability of evapotranspiration in the Rietholzbach basin determined using energy balance methods" at ETH Zrich, which was a continuation of his studies in the area of Natural sciences at the same university. Since his completion of the doctor studies he has lead research at ETH Zrich, the Swiss National Hydrological and Hydrometeorological Services, the Swiss Federal Offic for Water and Geology and the Swiss Academy of Natural Sciences. Since 2009 Dr. Schdler has joined the University of Bern as a senior researcher and lecturer on Climate Change and Hydrology. Previously he has been lecturing at ETH Zrich, Moscow University, University of Geneva, EPF Lausanne, Institut Universitaire Kurt Bschr.

Benefits of wireless mesh networks for environmental sensing within the research project MontanAqua

Emmanuel Rey, Bruno Schädler, Jan Schwanbeck and
Rolf Weingartner
OCCR and GIUB, University of Bern
Bern, Switzerland
jan.schwanbeck@giub.unibe.ch

Markus Anwander, Torsten Braun, Almerima
Jamakovic and Thomas Staub
IAM, University of Bern
Bern, Switzerland
anwander|braun|jamakovic|staub@iam.unibe.ch

Abstract—Continuously measured environmental variables are invaluable for research projects as MontanAqua. The extensive sensor network of MontanAqua was enhanced with the A⁴-Mesh wireless mesh network which leads to various benefits e.g. lower maintenance costs and reduced risk for data loss.

Keywords: *wireless mesh networks; environmental sensors, MontanAqua, A⁴-Mesh, authorization and authentication infrastructure (AAI)*

I. INTRODUCTION

Dry valleys in the Alps are particularly vulnerable to changes in water availability and water demand. MontanAqua, a transdisciplinary project of the national research programme NRP61 on sustainable water management, aims at developing water management strategies by studying in depth the Crans-Montana-Sierre region in Valais, Switzerland [1].

Extensive field-measurements and modelling of both water resources as well as water use will display the current state of the water balance. The knowledge of the current state together with likely scenarios of socio-economic and climate change offer the opportunity to develop strategies for sustainable and integral water management.

Technical advances in recent decades lead to availability of a wide range of reasonable priced sensors for hydro-meteorological observations. Costs for data-transfer and maintenance are factors that usually limit the number of observation plots, especially for sensors placed in very remote and hardly accessible sites.

In this paper we present the environmental sensor network of the MontanAqua project and the benefits drawn from its enhancement with A⁴-Mesh technology. A⁴-Mesh is a novel wireless mesh architecture that provides an important basis for the real-life deployment of wireless mesh networks for the support of environmental research. The main goal of the wireless mesh architecture is to enable easy and secure broadband network access at remote locations.

II. STUDY AREA

The study area (Fig. 1) extends over the Crans-Montana-Sierre region on the south facing slope of Bernese Alps. Boundaries are defined by the rivers Liène, Raspille and Rhone and the mountain ridge. Further the area is characterized by complex topography with altitudes ranging between 500 m

a.s.l. at river Rhone and 3000 m a.s.l. at the mountain ridge, strong precipitation and temperature gradients and a wide range of land cover types.

The following land covers are dominant along the main altitudinal gradient from the lowest slopes up to the mountain ridge: vineyards, intensive and extensive grassland and expanding mixed forests, tourist resorts and typical alpine landscape with coniferous forests, seasonal alpine pasture and rock as well as glaciers in highest zone.

III. ENVIRONMENTAL SENSORS

In 2010, as a part of the MontanAqua project, an extensive monitoring network (Fig. 1) has been set up for estimating and modeling the water availability under present and future conditions.




Figure 1. Study area and measuring network [2]

This network comprises twelve automatic weather stations, eighteen rain gauges, eight soil moisture plots, one hill-slope lysimeter, six runoff gauges and a high resolution webcam.

The automatic weather stations measure rainfall, temperature, relative humidity, solar radiation and in part wind velocity and wind direction. Most of them transmit data once a day via GSM/GPRS. Data from rain gauges, soil moisture plots, lysimeter and runoff gauges could initially only be read-out onsite. After deploying the A⁴-Mesh wireless mesh network the most remote weather stations and runoff gauges are accessible directly from campus network.

IV. WIRELESS MESH NETWORK (WMN)

Wireless mesh networks (WMNs) have been subject to intensive research for several years [3]. There is a diverse range of possible application scenarios for the deployment of wireless mesh networks. First outdoor applications of WMNs have been reported e.g. serving for local flood warning [4] or facilitating sensor data collection for ecological researchers in a Natural Reserve [5].

A⁴-Mesh is a novel wireless mesh architecture that supports secure broadband network access for researchers in remote areas. In order to do that, authentication and authorization mechanism are designed for end users as well as mesh nodes, giving the possibility of a seamless WMN integration into the home's organization authentication and authorisation infrastructure. This mechanism is designed upon existing federated access control approach, i.e. the AAI infrastructure that is using just a single user name and password among organization.

The current setup of the WMN (Fig. 2) consists of seven wireless mesh nodes interconnecting various hydro-meteorological sensors to the university campus network. The A⁴-Mesh WMN is a common form of WMNs where every mesh node relay data for other mesh nodes (a typical ad-hoc networking paradigm) and given mesh routers also have the additional capability of being Internet gateways. Node 1 serves as a gateway to the fixed glass fiber network backbone of SWITCH, the Swiss National Research and Education Network (NREN) operator. In this way it carries the traffic between the mesh nodes and the Internet. The first link from node 1 is directed to a relay station (node 2) in Vercorin at the opposite hill slope, where a high resolution webcam is located. The second link from node 1 connects to nodes 4b and 4a on Cry d'Er, which in turn interlink with all the other nodes except node 3.

Exposition to strong weather conditions in alpine environment with very low temperatures and a lot of snow in winter as well as thunderstorms with heavy rainfall, wind or hail in summer requests an appropriate design of A⁴-Mesh nodes: i.e. durable steel cases, self-contained power supply based on a solar panel at least three meters above ground to prevent them from being covered by snow and batteries powerful enough to bridge periods with poor sunlight. Additionally, as a way to improve lifetime and reduce maintenance cost of WMNs, the mesh nodes that are covered by the cellular networks have a backup mesh node with UMTS for remotely solving e.g. software issues.




Figure 2. Wireless mesh testbed with nine nodes and interlinks of up to almost 10 km serving the study area of the MontanAqua project in Valais, Switzerland with broadband network access.

V. BENEFITS OF WMNs FOR ENVIRONMENTAL SENSING

The environmental research project MontanAqua maintains an extensive hydro-meteorological monitoring network. The extensions of this pure measurement network with A⁴-Mesh technology for data transfer brought the following benefits:

- easy and cellular network independent data transfer from field sites to campus network at short intervals to prevent data loss,
- real time access to field sensors from desk at university for remote control in case of errors,
- low-cost broadband network access at field sites.

ACKNOWLEDGMENT

The A⁴-Mesh project is carried out as part of the program "AAA/SWITCH - e-Infrastructure for e-Science" lead by SWITCH, the Swiss National Research and Education Network, and is supported by funds from the Swiss State Secretariat for Education and Research.

REFERENCES

- [1] R. Weingartner, K. Herweg, H. Liniger, S. Rist, B. Schädler, F. Schneider, O. Graefe, M. Hoelzle, and E. Reynard, "Water Scarcity in Inner-Alpine Regions. Options for sustainable water use in Crans-Montana-Sierre region (Valais)", MRI News, No. 5, December 2010
- [2] E. Rey, M. Kauzlaric, R. Weingartner, B. Schädler, H. Liniger, and K. Herweg, "Approaching water stress in the alps – water management options in the Crans-Montana-Sierre region, Valais, Switzerland (MontanAqua)," Proceeding of SHF Conference "Water in mountain areas. Better observations for better forecasting", Lyon, mars 2011.
- [3] I.F. Akyildiz, X. Wang, and W. Wang, "Wireless mesh networks: a survey," Computer Networks Journal, vol. 47, pp. 445-487, 2005.
- [4] P. J. Smith, D. Hughes, K. J. Beven, P. Cross, W. Tych, G. Coulson, and G. Blair, "Towards the provision of site specific flood warnings using wireless sensor networks," Met. Apps, vol. 16, no. 1, pp. 57-64, 2009.
- [5] D. Wu, D. Gupta, and P. Mohapatra, "QuRiNet: A wide-area wireless mesh testbed for research and experimental evaluations," AD HOC NETWORKS, vol. 9, pp. 1221-1237, 2011.

10 Real-time spatiotemporal combination of radar and raingauge precipitation measurements in Switzerland

Ioannis Sideris

MeteoSwiss

E-mail: Ioannis.Sideris@meteoswiss.ch

Dr. Ioannis Sideris was born in Greece and got his bachelor's in Physics from the Aristotelion University of Thessaloniki. He continued his studies in USA where he earned his Master's and PhD in Astrophysics. His main research has been on chaos on galactic systems, stochastic processes, and the N-body problem. He continued his career in the Northern Illinois University where he was involved in research on beam dynamics and accelerator physics in association with Fermi Lab. Then, he moved to the Department of Computational Astrophysics in the University of Zurich. Currently he is a scientific collaborator in MeteoSwiss Locarno, where he works with the radar team on research and operational implementation of applications related to quantitative precipitation estimation. He has authored a number of scientific papers and has given several talks on scientific topics of a variety of fields.

Real-time spatiotemporal combination of radar and raingauge precipitation measurements in Switzerland.

Sideris Ioannis¹, Gabella Marco¹, and Germann Urs¹

¹MeteoSwiss, Stabile La Ferriera, Via Luini 12, CH-6600 Locarno
(Ioannis.Sideris@meteoswiss.ch)

The main needs for precipitation estimation in Switzerland is served by two types of monitoring devices: a dense raingauge-network and three C-band radars. Raingauges and radars operate independently and each is characterized by well-known advantages and disadvantages. Raingauge measurements are considered to be fairly accurate and provide direct precipitation on the ground, however their representativity is questionable: in spatial aspects they are equivalent to point measurements, and the rainfall information they provide depends strongly on how densely populated the available raingauge-network is and how complex the meteorological situation is. For instance, when the weather is characterized by strong convective cells, statistical microvariability may largely prohibit raingauges from sufficiently capturing the profile of the precipitation process. Moreover, strong winds, exceptionally high rain-rate, and snow may affect to a certain degree the accuracy of raingauge measurements. On the other hand radars are characterized by dense and widespread estimation of rainfall. The typical spatiotemporal resolution of a radar is about 1km² every 5 minutes. This monitoring however is entangled with significant errors, for example those associated with physical obstacles, such as the orographic ones, or those related to choice of mathematical factors to be used in translating the backscattered radiation (which radars actually measure) into rain-rate. The main point is that radars monitor precipitation densely but indirectly, and in a terrain as challenging as the swiss alpine one, errors naturally emerge.

Although the spatiotemporal resolution of the radar image is excellent, one naturally wonders if we can incorporate both radar and raingauges into a common scheme to improve maps of precipitation estimates. Researchers started attacking this problem since the eighties employing a number of schemes [1-9]. Prominent among them stand the geostatistical techniques which aim in zeroing the bias of the error between the estimator and the actual value of precipitation but simultaneously minimizing the variance of this error. The technique which is commonly used in the context of quantitative precipitation estimation (QPE) is coined as “kriging with external drift” (KED). This describes a statistical interpolation method where precipitation is taken as a stochastic process and a trend (or drift) is determined with the help of the radar estimation map. The expectation is that the raingauge-radar merging such technique implies is able to provide an improved (in terms of errors) picture of rainfall estimates.

Although merging schemes have been applied and discussed for targeted events within the context of academic research, only in the last decade they started emerging as operational tools for providing real-time rainfall estimates. The task is not trivial: real-time operational environment assumes no-human-intervention, meaning that an application should be based on a scheme sufficiently robust to deal with delicate and practical matters, such as complex meteorological conditions, or errors of significant size.

Over the last two and a half years an effort initiated within MeteoSwiss (NCCR III-CombiPrecip) to provide raingauge-adjusted rainfall maps. This was motivated in part by sizeable losses suffered due to severe floods which took place in 2005 and 2007. Aiming in utmost robustness and stability of the merging application we expanded the idea of the aforementioned geostatistical scheme. Kriging with external drift is spatially based; we succeeded introducing also temporal information into this scheme, in the form of additional variables. The scheme that will be incorporated is now co-kriging with external drift (CED) with the term “co-” signifying temporal information.

The performance of our application will be discussed in terms of (a) accuracy, (b) stability, (c) speed, and (d) flexibility. Regarding accuracy several skill scores will demonstrate how our scheme succeeds in improving rainfall estimates in comparison to radar. Moreover, pictures and animations or rainfall maps will be provided so the audience gets a visual impression of how the raingauge-radar merging operates in actual terms.

In terms of stability it will be shown how CED, due to its spatiotemporal structure, succeeds providing reasonable estimates even in cases where robustness of input data is doubtful. It will be shown that in such cases KED often fails to provide realistic estimates, a fact that in an automatized operational environment would render the tool unstable and arise questions regarding the reliability of the product.

In terms of speed it will be shown that the merging can be achieved in less than 10 minutes using the computational resources of MeteoSwiss. This gives us the opportunity to set a plan regarding the optimal output of this application. At this stage the expectation is to produce hourly-aggregated precipitation maps of Switzerland every 10 minutes, with the hourly aggregations being motivated by the necessity for input sufficiently robust.

In terms of flexibility, it will be discussed that needs of potential clients vary, many times significantly, a factor which is central to be taken into consideration during development. For instance the fact that our application will produce hourly maps only in the area of Switzerland may make it impractical for nowcasting needs. Moreover, geostatistical interpolation (as is commonly true for other interpolation schemes) gradually fails as distance increases from the monitored terrain (in our case out of swiss borders). These are practical issues important to resolve. A disaggregation scheme will be presented which divides proportionally the hourly precipitation maps into twelve five-minute precipitation maps, thus producing a raster spatiotemporally equivalent to that of the radar, but having incorporated the geostatistical adjustment from the raingauge-radar merging. Finally, the entirely novel concept of “virtual raingauges” will be presented to demonstrate how it improves the situation on extrapolation issues and how it might help with microvariability problems common in convective cases.

References

- [1] J. Creutin, G. Dekrieu, T. Lebel, *J. of Atmosph. and Ocean. Technology*, **5(1)** (1988), 102.
- [2] R. Erdin, M.S. Thesis, Mathematics, ETH Zurich (2009).
- [3] M. Gabella, J. Joss, G. Perona, *Journal of Geophys. Research*, **105** (2000), 2237.
- [4] M. Gabella, R. Notarpietro, *IEEE Geoscience and Remote Sensing Letters*, **1** (2004), 1545.
- [5] U. Germann, G. Galli, M. Boscacchi, M. Bolliger, Q. J. Roy. Meteor. Soc., **132** (2006), 1669.
- [6] W.F. Krajewski, *Journal of Geophysical Research*, **92** (1987), 9571.
- [7] M. Salek, P. Novak, Proceedings of ERAD, Zurich, **251** (2008).
- [8] D. Seo, W.F. Krajewski, D.S. Bowles, *Water Resources Research*, **26** (1990), 469.
- [9] C.A. Velasco-Forero, D. Sempere-Torres, E.F. Cassiraga, J.J. Gomez-Hernandez, *Advances in Water Resources*, **32** (2009), 986.