

EWSN 2014: Posters and Demos

Welcome to Oxford, U.K., and to the Poster and Demo Session of the 11th European Conference on Wireless Sensor Networks (EWSN).

EWSN is a highly selective conference in the area of sensor networks and the Poster and Demo session provides an ideal platform to disseminate early research results and to showcase exciting demonstrations. In addition to the traditional Poster and Demo session, the authors of accepted submissions participate in a 1-minute madness session where they give a short pitch to all EWSN attendees.

This year we received a total of 43 submissions from authors in 20 different countries representing all continents, except Antarctica! Considering that poster and demos are aimed at disseminating preliminary work, the main criteria for acceptance was novelty. After an initial evaluation, submissions were divided into three categories: accept, borderline and reject. Most of the borderline submissions were discussed in detail among the three co-chairs. We accepted 29 submissions, 13 demos and 16 posters. The accepted posters and demos cover a wide range of topics, from traditional networking research such as routing, time synchronization and data fusion, to smart environments and novel localization methods.

The Poster and Demo session is an integral part of the EWSN conference and we want to express our gratitude to the General Chair, Niki Trigoni, the TPC Co-Chairs, Amy Murphy and Bhaskar Krishnamachari, and the members of the Steering Committee for their constant support.

We hope that you will enjoy the program and that you will strongly consider submitting poster and demos to future versions of EWSN.

Poster and Demo Co-Chairs EWSN 2014

Alberto Cerpa, University of California, La Merced, USA

Timothy Hnat, University of Memphis, USA

Marco Zuniga, Delft University of Technology, The Netherlands

Contents

Applications	4
Demo Abstract: Data Center Monitoring on a Budget <i>Peter Haban, Faye Mitchell, and Cigdem Sengul</i>	4
Demo Abstract: Lessons Learned from a WSN Deployment for Precision Agriculture <i>Prashant A. Patil, Jabal Raval, Bhushan G. Jagyasi, Nilima Warke, and P. P. Vaidya</i>	6
Demo Abstract: Measuring Human Energy Expenditure with K-Sense <i>Kazi I. Zaman, Anthony White, Sami Yli-Piipari, and Timothy W. Hnat</i>	8
Demo Abstract: Demo Abstract: Object Recognition in Visual Sensor Networks. A comparison of two different operative paradigms <i>A. Canclini, L. Baroffio, M. Cesana, A. Redondi, and M. Tagliasacchi</i>	10
Poster Abstract: Sensor Network Based Surveillance System for Forest Protection <i>Gergely Vakulya, Gyula Simon, László Czúni, and Péter Z. Varga</i>	12
Localization and Tracking	14
Demo Abstract: Lightweight Continuous Indoor Tracking System <i>Zhuoling Xiao, Hongkai Wen, Andrew Markham, and Niki Trigoni</i>	14
Demo Abstract: Multi-Sensor Fusion for Indoor Pedestrian Tracking <i>Jan Medvesek, Andrew Symington, and Stephen Hailes</i>	16
Poster Abstract: A Case for Magneto-Inductive Indoor Localization <i>Traian E. Abrudan, Andrew Markham, and Niki Trigoni</i>	18
Poster Abstract: WiFi Sensors Meet Visual Tracking For An Accurate Positioning System <i>Savvas Papaioannou, Hongkai Wen, Zhuoling Xiao, Andrew Markham, and Niki Trigoni</i>	20
Poster Abstract: Robust Phase Estimation for RSSI signals <i>Gergely Zachár and Gyula Simon</i>	22
Poster Abstract: Robust Sensor Fusion for Radio Interferometric Tracking <i>Gergely Zachár and Gyula Simon</i>	24
Smart Buildings	26
Poster Abstract: Delivering Intelligent Home Energy Management with Autonomous Agents <i>Conor Muldoon, Michael J. OGrady, David Lillis, Tadhg OSullivan, Thomas Holz, and Gregory M.P. OHare</i>	26
Poster Abstract: Experimental Analysis of Environmental Perturbations on Wireless Sensor Network Operation <i>Mustafa Mousa, Ahmad Dehwah, and Christian Claudel</i>	28
Poster Abstract: Intelligent Building Management Systems Based on Cognitive M2M Communications <i>Manolis Surligas, Stefanos Papadakis, Vangelis Angelakis, Sofoklis Kyriazakos, Elias Tragos, and Di Yuan</i>	30
Poster Abstract: Secure Metering Communication in the Energy Smart Home Lab <i>Anton Hergenröder, Kaibin Bao, and Thomas Bräuchle</i>	32

Energy	34
Demo Abstract: Activity Rate Adaptation in WSN	
<i>Victor Cionca, Wensi Wang, and Brendan OFlynn</i>	34
Poster Abstract: Sustainable Backpressure Data Collection in Solar Powered Wireless Sensor Networks	
<i>Shusen Yang and Julie A. McCann</i>	36
Testbeds and the Cloud	38
Demo Abstract: SicsthSense - Dispersing the Cloud	
<i>Liam McNamara, Beshr Al Nahas, Simon Duquenooy, Joakim Eriksson, and Thiemo Voigt</i> .	38
Demo Abstract: How Temperature Affects IoT Communication	
<i>James Brown, Utz Roedig, Carlo Alberto Boano, Kay Römer, and Nicolas Tsiftes</i>	40
Demo Abstract: Testbed Infrastructure for Benchmarking RF-based Indoor Localization Solutions under Controlled Interference	
<i>Filip Lemic, Jasper Büsch, Mikolaj Chwalisz, Vlado Handziski, and Adam Wolisz</i>	42
Protocols	45
Demo Abstract: Foren6, a RPL/6LoWPAN Diagnosis Tool	
<i>Sébastien Dawans and Laurent Deru</i>	45
Demo Abstract: MPASS: Multi-channel Packet Sniffing System based on IEEE 802.15.4	
<i>Jung-Hwan Bae, Tae-Soo Kim, Hiecheol Kim, and Seong-eun Yoo</i>	47
Demo Abstract: Reliable 6LoWPAN Networks with Multiple Border Routers	
<i>Laurent Deru and Sébastien Dawans</i>	49
Poster Abstract: Decentralized Time-synchronized Channel Swapping for Wireless Sensor Networks	
<i>G. Smart, N. Deligiannis, Y. Andreopoulos, R. Surace, V. Loscri, and G. Fortino</i>	51
Poster Abstract: Fairness in WSNs	
<i>Yad Tahir and Julie A. McCann</i>	53
Poster Abstract: AGRO - Optimal Routing with Low Latency and Energy Consumption in Wireless Sensor Networks with Aerial Relay Nodes	
<i>Jin Zhang, Kai Li, Andreas Reinhardt, and Salil S. Kanhere</i>	55
Poster Abstract: Neighborhood Cardinality Estimation in Extreme Wireless Sensor Networks	
<i>Marco Cattani, Marco Zuniga, Andreas Loukas, and Koen Langendoen</i>	57
Poster Abstract: Synchronizing Trickle Intervals	
<i>Joakim Eriksson and Omprakash Gnawali</i>	59
Poster Abstract: The Purpose and Benefit of Segmenting a Multi-Transmitter Network	
<i>Douglas Carlson and Omprakash Gnawali</i>	61

Demo Abstract: Data Center Monitoring on a Budget

Peter Haban and Faye Mitchell and Cigdem Sengul
Computing and Communication Technologies
Oxford Brookes University
Email: {phaban,frmittchell,csengul}@brookes.ac.uk

Abstract—In this demo, we will demonstrate a system that monitors the power usage of a small scale data center. Our system is installed in the Oxford Brookes University data center as a proof of concept. The design offers the necessary flexibility to support the expected complexities of a small organically growing data center.

I. INTRODUCTION

Several studies have shown that 1.3% of world wide electricity use is attributed to data centers [1] and 30% of the energy cost of a data center is attributed to cooling [2]. On the positive side, recent equipment allows for higher temperatures on the air intake [3] and raising the intake temperature has been shown to be effective in saving energy [4], [5]. However, reliable monitoring and alerting mechanisms are required to ensure temperature thresholds are not exceeded to avoid component failure and resulting downtime. While established mechanisms are available for monitoring and optimising power usage in large data centers, they are typically out of reach for smaller data centers. Hence, the aim of this research is to fill this gap and provide an economic monitoring solution.

Current research suggests that an integrated system of wired and wireless sensors may provide environmental measurements at very high temporal and spatial resolution [6]. However, proposed solutions are typically aimed at large-scale data centres. In this demo, we will show a light-weight, low-cost, and efficient solution based on open hardware and software components. In addition, we have implemented an alarm monitors, and visualization support. Our proposed system has been developed specifically to fit a wide range of scenarios including small scale deployments at reasonable cost.

II. DESIGN OF A MONITORING SYSTEM

A. System Overview

Hardware. Our system is composed of a server and sensor nodes that have multiple temperature sensors. The server implements monitoring, alerting, and visualization services, which are described in Section II-B. It also implements a database for storing temperature readings. The sensor node platform is Raspberry Pi. We use Power over Ethernet (PoE) to power the nodes and communicate sensor readings. The support for Raspberry Pi needed to be developed as it did not exist at the time of our development. For temperature sensors, we use OneWire sensors. In addition to their low cabling requirements, these sensors support a common bus

arrangement which allows to scan for all sensors on the bus and directly access individual sensors with their specific address. Based on the initial readings, it was found sufficient to record readings in one minute intervals to build heat maps. Even though, we use wired communication for our proof-of-concept deployment in Section III, our system can work on any type of communication, wired and wireless. In the demo, we will demonstrate its wireless operation.

Software. Both the server and sensor node comprise of a full Debian LAMP stack and additional logic implemented in Python executed via Cron (i.e., services described in Section II-B). We chose to use well-used open source software when implementing our services, rather than, for instance, using MQTT, CoAP, or XMPP, as we aimed to provide a solution with low-entry barrier for its immediate users. Unique to the server side are reporting and visualization elements based on HTML Canvas. Elements unique to the sensor node side are the PHP Slim framework which provides the RESTful web services and the PHP logic which executes CLI commands.

B. Supported Services

The monitoring of the data center is accomplished through (1) discovery and collector service, which also includes an alerting mechanism and (2) visualization services, which are detailed next.

1) *Discovery and Collector Service:* We implemented two alternatives for this service. In *distributed mode*, each sensor node is pre-configured with the server address. The local collection service on each node accesses all local sensors and uploads readings in set intervals to the central server database. In contrast, in *centralized mode*, each node offers three RESTful web services which allow identifying the nodes, the sensors on each node, and readings from each sensor. To this end, the server-side discovery service scans a range of addresses. Identified nodes and their sensors are stored in the database. Also, using a collection service, the server periodically harvests readings from all registered sensors.

We have experimented with both distributed and centralized modes, and chose to work with the centralized mode in our deployment due to not having to preconfigure nodes. Furthermore, any change to the server could easily be handled without requiring a reconfiguration of all nodes. It also becomes possible to scale the system to a large number of

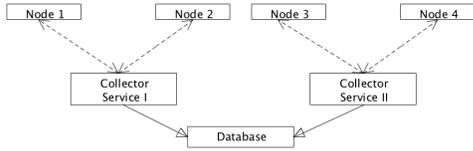


Fig. 1. Centralized mode provides scalability through additional collectors.

nodes through provisioning of additional collector services as shown in Figure 1.

Finally, we also implemented an alerting mechanism as part of the collection service. Before uploading readings to the database, the collection service compares each reading to a pre-defined threshold, and sends alert e-mails when temperature readings exceed the threshold. The threshold and the credentials for the alert e-mail account are configured through the server-side management interface.

2) *Visualization*: We built several interfaces to visualise the sensor readings. These include physical layout of the room and temperature readings for the selected point in time as well as a hyperlink over each component stack to link to the management interface of the equipment.

III. PROOF-OF-CONCEPT DEPLOYMENT

We deployed our system in a small data center at Oxford Brookes University, which consists of two double rows of 19" racks with cold isle containment and a separate room which houses the UPS battery stack. Cold air is supplied by several cooling units which are located around the cold isles and feed cold air into the under floor cavity. Floor grids in the cold isles and UPS room allow to direct cold air. The longer cold isle is currently 50% populated, partially with networking equipment. The short cold isle is about 70% populated including a large tape library.

We fitted four Raspberry Pis including cable runs along both sides of the cold isles which will also allow additional sensors. Eight sensors have been installed to get readings from the cold isles (sensors 1 and 2 of nodes 1-4 in Fig. 2). Also, two sensors are monitoring air intake temperature of the cooling units (sensors 2/3 and 3/3), one sensor is monitoring the long cold isle (sensor 1/3), and one sensor is monitoring the temperature in the UPS room (sensor 4/3).

The average recorded temperatures during our evaluation are presented in Table I. The sensors of node 4 all report low temperatures, which is for the UPS room. This room has a target temperature of 20°C. The sensor 3 of nodes 1, 2 and 3 report higher temperatures as they are not located in the cold isles. Hence, the temperature readings represent the current data-center set-up. As a next step, we plan to run an audit to determine the upper temperature threshold which will be based on the lowest required air-intake temperature. We believe that with the developed monitoring and alerting facilities, it will be possible to safely raise the base temperature.

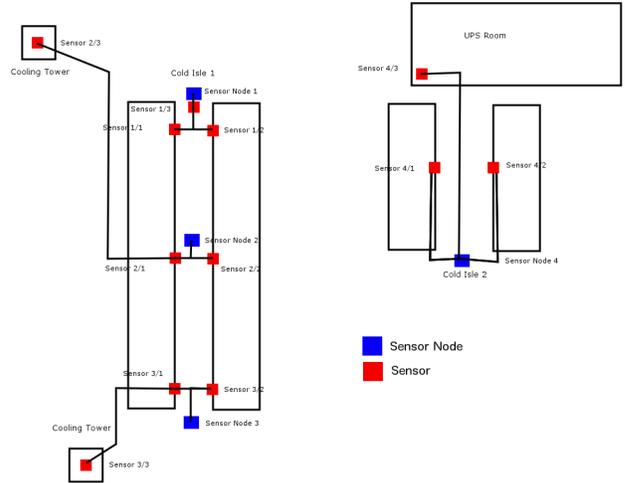


Fig. 2. Deployment in a small-scale data center.

TABLE I
TEMPERATURE READINGS

Sensor node	Sensor 1	Sensor 2	Sensor 3
I	19	19	23
II	15	16	23
III	14	18	22
IV	16	16	18

IV. DEMO SCENARIO

The demo will showcase the main components of our monitoring solution and demonstrate how sensor readings can be collected from the nodes. The demo setup will consist of two wireless sensor nodes, a wifi router and a laptop acting as the server. This will include a demonstration of the configuration interface and Web Services offered by each sensor node as well as the server side configuration and reporting interfaces.

V. CONCLUSION

In this demo, we present an open-source solution which can create a light-weight sensor network at a minimal cost, avoiding the risk of vendor lock-in for a small-scale data center. By combining mobile and web development technologies, and highly integrated sensors, it also becomes possible to scale our solution as the data center organically grows.

REFERENCES

- [1] J. G. Koomey, "Worldwide electricity used in data centers," *Environmental Research Letters*, 2007.
- [2] N. Rasmussen, "Calculating total cooling requirements for data centers," American Power Conversion, Tech. Rep., 2007.
- [3] *Site Planning Guide for Sun Servers.*, Oracle, 2007.
- [4] M. Patterson, "The effect of data center temperature on energy efficiency," in *11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems ITherm*, 2008, pp. 1167–1174.
- [5] M. R., "Google: Raise your data center temperature," Tech. Rep., Apr. 2013.
- [6] Pereira Nuno and Stefano Tennina and Eduardo Tovar, "Building a microscope for the data center," *Wireless Algorithms, Systems, and Applications*, 2012, pp. 619–630.

Demo Abstract: Lessons Learned from a WSN Deployment for Precision Agriculture

Prashant A. Patil*[†], Jabal Raval[†], Bhushan G. Jagyasi[†], Nilima Warke* and P. P. Vaidya*

[†] TCS Innovation Labs Mumbai and * V.E.S.I.T. Chembur

{prashant.patil2, jabal.raval, bhushan.jagyasi}@tcs.com, {nilimavwarke, vaidya0612}@gmail.com

Abstract—The paper present some key learnings from a design and deployment of wireless sensor networks for precision agriculture applications. We have used solar powered IRIS motes with externally interfaced sensors. The field parameters recorded by sensor nodes have been used to forecast the late blight disease risk in the potato crop. We believe that the insights gained from the learning will enable the robust WSN deployment for outdoor environments.

I. INTRODUCTION

Wireless Sensor Networks (WSN) are formed by a large number of sensor nodes deployed in an application area with sensing, processing and communication capabilities. The Wireless Sensor Networks have been shown to be useful for a wide range of applications from battlefield, surveillance, disaster detection to the precision agriculture. In mKRISHI® project [1], the data from WSNs is used for providing the agro advisory services to the rural farmers. In [2], we proposed the utility of WSNs in the Potato late blight disease forecasting system.

In this paper, we present the design of a low cost sensor node based on the IRIS platform [4] and also demonstrate the key learnings from its deployment in the Merino Potato farms at Hapur in Uttar Pradesh (UP), India.

The WSNs have been deployed in the potato farms for monitoring the canopy temperature and canopy humidity parameters in order to provide Late Blight Disease alerts based on Central Potato Research Institutes’s Jhulsacast model [3] to the potato farmers of UP, India. While designing the enclosure and the interfacing circuits, we considered various practical constraints based on our prior experience on agricultural parameter monitoring. Once deployed, WSNs are expected to sustain and supply data uninterruptedly for the full crop cycle which may vary from a few months to a year. Further, the external environment needs the sensor nodes to be more rugged. Available pro version sensor nodes [4] and environment monitoring systems for precision agriculture are very expensive to be suitable for providing services to the farmers. We believe that our approach of the low cost nodes without compromising the accuracy and reliability of the sensor data has made the design more appealing.

II. SENSOR NODE DESIGN

In this section, we present the design of sensor node using the IRIS mote (XM2110) and sensor board (MDA300CA) for monitoring canopy temperature and canopy humidity from the

agricultural farms. A signal conditioning board has been designed for making signals from interfaced sensors compatible with the specifications of the IRIS mote. Whole set up is assembled in waterproof enclosure as shown in figure 1.



Figure 1: IRIS based Sensor Node

A. Signal Conditioning Board

As shown in figure 2, voltage divider is used to scale down the output of humidity sensor in 0 to 2.5 V range as per the input specifications of ADC on MDA300CA. Maximum output of temperature sensor is already below 2.5V so signal conditioning is not required for interfacing this sensor to ADC.

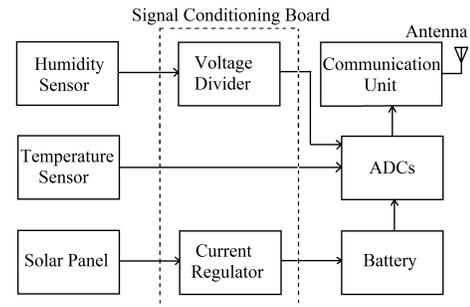


Figure 2 : Block diagram of a sensor node

The circuit for harnessing solar power to charge the batteries is also included on the signal conditioning board. We are using LM317T chip for regulating current output of solar panel.

B. Sensor Modules

We have used easily available sensors like LM35 for temperature and SY-HS-230 for humidity. Unshielded multi core wire and 3.5 mm audio jack are used for plugging sensors to a node.

C. Data gathering from WSN

The field data reaches local machine through a base node. Local machine interprets the hexadecimal data and stores it

in local database. We wrote Java code on Base station to communicate all received data to the servers. Once the data is received, the server side code converts them into actual parameters. We chose back end calibration of the nodes as it provides more flexibility to change the equations of conversion if necessary for calibration. Scheduler has been deployed in order to send data at fixed intervals to the servers. For the 12 bit ADC, the maximum number of divisions possible are 4096 and considering 2.5 V as reference for ADC, we compute temperature T in $^{\circ}\text{C}$ as $T = 2.5A \times 100/4096$, where A is the hexadecimal equivalent of temperature and scaling factor of 100 is used as LM35 provides 10mV output per $^{\circ}\text{C}$ rise in temperature.

Non-linearity of the humidity sensor output has compelled us to use different conversion equations (refer Equation(1)) for the different range of humidity.

$$H = \begin{cases} -1.828B + 2.052 \times 10^{-3}B^2, & 900 < B < 930 \\ -3.141B + 3.456 \times 10^{-3}B^2, & 930 < B < 940 \end{cases} \quad (1)$$

where B denotes the value of humidity in hexadecimal and H represents % Humidity. We have similar equations for other ranges of B .

III. FIELD DEPLOYMENT RESULTS

The Figure 3 represents the deployment of the Wireless Sensor Network in potato farms at Hapur UP. Here sensor node 1 act as a parent and sensor node 2 as a child. We have equipped parent node with 3V 2700 mAh while child node have 3V 2000 mAh batteries. This is to improve power management in a multi-hop scenario which in turn helps WSN to sustain for the longer durations.

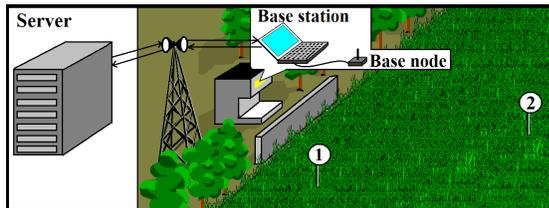


Figure 3 : WSN in potato field

Figure 4 presents the humidity and temperature observations, for a small representative window, from sensor node 2 in comparison to the actual readings measured manually from a standard digital sensing device. Figure 5 presents humidity and temperature readings of 4 consecutive days from sensor node 2.

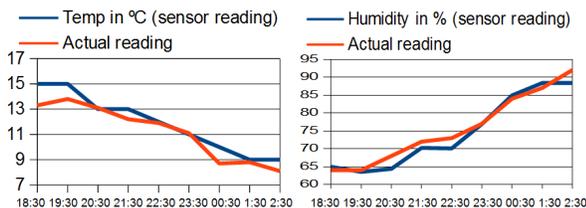


Figure 4 : Sensor data with respect to actual observation

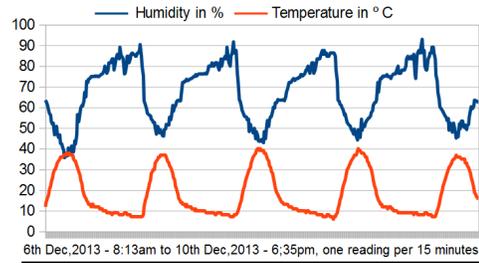


Figure 5 : Temperature and Humidity vs. Time for Node 2

Figure 6 presents hourly battery voltage of sensor node 2 for 24 hours. It can be observed that battery voltage never fall below 2.4V, minimum voltage required for node to work.

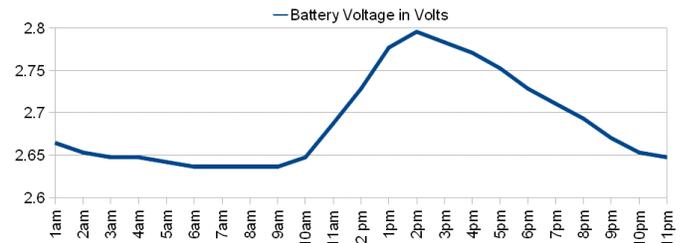


Figure 6 : Battery Voltage vs. Time for Node 2

The battery which we have used has the capacity of 2000mAh and the 3V solar panel provides current up to 300mAh. The sensor node uses 17mAh for the communication, computing and sensing purposes. India receives on an average 7 hrs sunlight daily. So batteries will charge to 99% (Approximately 1980mAh) of its full capacity, daily, if started from the zero charge. Sensor node consume 14% (Approximately 289mAh) battery in remaining 17 hrs. At start of next charging cycle batteries will still have 85% (Approximately 1690mAh) of charge which means sensor node can work for next 4 days without sun light.

The cost of IRIS based node is found to be one third as compared to the cost of eKo node [4]. Further the observations obtained from the deployed sensors are seen to be reasonably accurate in order to cater to the agro advisory and disease forecasting services to the farmers.

ACKNOWLEDGMENT

We are thankful to the Merino Farms at Hapur Uttar Pradesh for supporting us in deployment of WSN.

REFERENCES

- [1] Ajay Mittal, Chetan K. P., Srinivasan Jayaraman and Bhushan G. Jagyasi, "mKRISHI Wireless Sensor Network Platform for Precision Agriculture", 6th International Conference on Sensing Technology, 2012.
- [2] A. Pande, B. G. Jagyasi, and R. Choudhuri, "Late blight forecast using mobile phone based agro advisory system", in Proc. 3rd International Conference on Pattern Recognition and Machine Intelligence, PReMI 2009, Dec 2009.
- [3] JHULSACAST Model, in CPRI Website, Dec 2013, cpric.ernet.in/Jhulsacast.htm, last accessed Dec 2013.
- [4] Eko pro Series and IRIS platform, in Memsic Website, July 2012, www.memsic.com, last accessed July 2012.

Demo Abstract: Measuring Human Energy Expenditure with K-Sense

Kazi I. Zaman, Anthony White, Sami Yli-Piipari, Timothy W. Hnat
University of Memphis, Memphis, TN, USA
{kizaman,arwhite5,srylppri,twhnat}@memphis.edu

Abstract—Accurate energy expenditure monitoring will be an essential part of medical diagnosis in the future, enabling individually-tailored just-in-time interventions. However, there are currently no real-time monitors that are practical for continuous daily use. In this demo, we will show a prototype of the *K-Sense* energy expenditure monitor that uses inertial measurement units (IMUs) mounted to an individual’s wrist, waist, and ankle to determine angular velocity and position. The system utilizes inverse kinematics to determine the amount of energy required for each limb to achieve its current movement.

I. INTRODUCTION

Obesity has reached epidemic proportions throughout the United States and was recently classified as a disease by the American Medical Association. The causes of obesity are as broad as the number of instances; however, most are contributed to individual lifestyles. Research has shown that a small amount of exercise is beneficial, but motivating individuals to change their behaviors is difficult. Most commercially available systems rely on waist or wrist mounted accelerometers which fail to capture the motion of body extremities. Technologies such as metabolic carts [3] and calorimeter rooms [7] allow researchers to accurately measure energy expenditure; however, they are difficult and intrusive to wear or place a significant burden on the participant.

In this demo, we will show *K-Sense* energy expenditure monitor: a more accurate wearable monitoring system based on inertial measurement units (IMUs) [8]. Its main challenge is correctly measuring *energy expenditure*: the amount of energy a human body uses performing activities. *K-Sense* considers the efficiency of limb motion and basal energy expenditure to determine energy expenditure. Our vision of the system involves integrating these techniques into a smartphone, wrist watch, and shoe, resulting in a solution without requiring any extraneous wearable devices. We will show a prototype where a person will perform various activities and energy expenditure and body position will be calculated simultaneously.

II. K-SENSE

The primary goal of *K-Sense* is to measure a person’s limb movements as they go about their daily lives. We briefly describe hardware design and operation, signal processing, and energy estimation techniques.

K-Sense sensors utilize a Sparkfun Razor 9DoF inertial measurement unit (IMU), consisting of three axis accelerometers (ADXL345), gyroscopes (IDG3200), and magnetometers

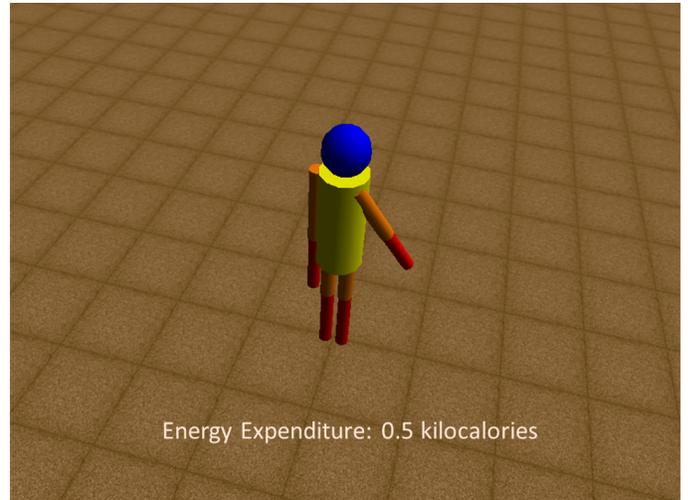


Fig. 1: A screen shot of *K-Sense*’s real-time visualization and energy expenditure monitor.

(HMC5883L), *K-Sense* sensors can be secured to a person’s wrist, ankle, and waist, as shown in Figure 2a by mounting on elastic bands. These three sensors are sampled by an on-board ATmega328 and send through a serial interface to our data collecting system. Firmware was modified to sample and output data as fast as 50Hz per sensor-axis or 450Hz for each IMU board, and transmitted via a Bluetooth serial interface (RN-41) at 115200 bits per second. A computer connects to all three devices over separate Bluetooth channels to log the data.

Data is collected by the IMU board and processed by our logging software where periodic errors in the time stamps occur. Errors appear as triangular shaped deviations where the logged time stamp is in the future. A robust linear regression technique and a predefined offset of 0.3 seconds is utilized to correct the errors in our data traces, effectively mapping the samples to a linear time sequence. Each IMU produces data with three degrees of freedom for each sensor type or nine degrees of freedom for the whole board. It transforms the corrected raw data streams into quaternions [4] which is a representation for the sum of a scalar and a three dimensional vector. Data is processed into 5 second windows to correlate with our ground truth energy expenditure device, which produces energy information every time a breath is taken.

We utilize angular sensor data, provided by the IMU, to

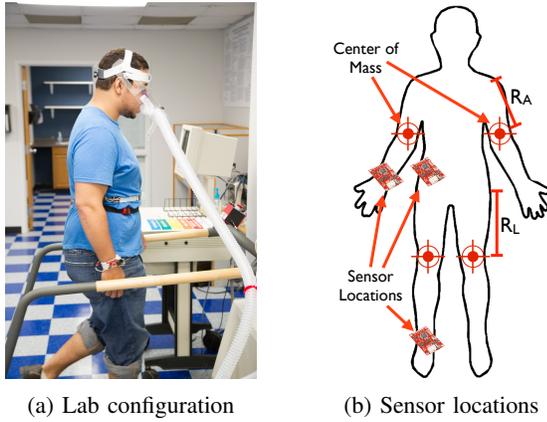


Fig. 2: (a) The K-Sense evaluation environment where a person is walking on a treadmill in a metabolic lab. (b) Sensors are placed at the wrist, ankle, and waist on the right side of each participant's body. The distance between a joint and a limb's center of mass, R_A and R_L , is used in the kinematic model.

estimate amount of work necessary to accomplish that movement. We place sensors on the right side of the body on the wrist, ankle, and waist as shown in Figure 2b. We are currently assume that each limb will expend equal amounts of energy whether the arms are moving as a mirror image to each other, or the arms and legs are moving in a counteracting manner and asymmetrical movement is not considered. We also assume that total energy expenditure is a function of the angular movement of the arms and legs along with the basal metabolic rate for the core of the body, the amount of energy expended to maintain the body's core functions. K-Sense estimates energy expenditure as follows where rotational work, $W = \tau\theta$, is a function of τ , rotational force and θ , angular displacement. By combining momentum, I , and angular acceleration, α , work becomes $W = I\alpha\theta$. Angular acceleration, $\alpha = \alpha_c/r$, is a function of tangential acceleration, α_c and the radius, r , from the limb joint to the limb center of mass (Figure 2b) resulting in work being defined as $W = I(\alpha_c/r)\theta$. Finally, by combining angular velocity, ω , with equations $\alpha_c = \omega^2r$, and $I = mr^2$ we derive

$$W = mr^2\omega^2\theta \quad (1)$$

where work, W , is a function of mass, m , the radius squared, the angular velocity squared, and the angular displacement. This equation forms the basis of K-Sense's energy estimator. The height and weight of a person are used in conjunction with equations derived from Paolo et al. [2] and Plagenhoef et al. [5] to compute the mass and lengths of individual body parts for participant and the weight is used to compute basal energy requirements. The basal energy consumption of the human body is approximately 4 kilojoule per kilogram of body weight per hour [1].

The following equation is used to estimate the total amount of energy expended by each participant

$$E_{estimate} = 2xE_{Wrist} + 2yE_{Ankle} + zE_{Body} \quad (2)$$

where the estimated energy is two times the energy measured at the wrists plus two times the energy measured at the ankles and the energy of the body. We have included a calibration factor, x, y, z , for each of the energy estimation components which are tuned by utilizing *minimax* optimization process to minimize root mean square (RMS) error or total energy error. K-Sense is able to achieve an accuracy of 92 percent with a standard deviation of 5 percent by utilizing the ankle, wrist, and body sensors with our kinematic model.

III. DEMONSTRATION

In this demonstration, we will illustrate how K-Sense measures human skeletal movement with its IMU sensors and through the application of inverse body kinematics. Energy expenditure is simultaneously computed and displayed. IMU sensors are placed around a person's waist, wrist, and ankle which will correspond in the future to a smartphone, smart watch, and smart shoe. We will demonstrate basic movements such as standing, sitting, walking, or jumping and show how the sensors process the data and produce energy estimates. An example of the visualization is shown in Figure 1.

For example, when an arm is moved, the wrist IMU measures accelerometer, gyroscopic, and magnetometer changes and transmits them at 50 hertz to a computer. We utilize inverse kinematics from Simbody [6] to compute the angular and position changes which can result in a determination of the arm's spatial and angular position. By utilizing inverse kinematics, we are able to estimate the amount of work necessary to achieve these positions and therefore compute the total amount of energy necessary. For a more complete description of K-Sense, please see our paper "K-Sense: Towards a Kinematic Approach for Measuring Human Energy Expenditure" [8].

REFERENCES

- [1] Fones,energy consumption and body weight. <http://www.bwl.admin.ch/themen/00509/00528/index.html?lang=en>. [Federal Department of Economic Affairs and Education and Research EAER; Online; accessed 4-September-2013].
- [2] P. de Leva. Adjustments to zatsiorsky-seluyanov's segment inertia parameters. *Journal of Biomechanics* 29 (9), pp. 1223-1230., 1996.
- [3] K. Jensen, S. Jorgensen, and L. Johansen. A metabolic cart for measurement of oxygen uptake during human exercise using inspiratory flow rate. *European Journal of Applied Physiology*, 2002.
- [4] S. Madgwick. An efficient orientation filter for inertial and inertial/magnetic sensor arrays. *Technical report Department of Mechanical Engineering University of Bristol*, april 2010.
- [5] S. Plagenhoef, F. G. Evans, and T. Abdelnour. Anatomical data for analyzing human motion. *Research Quarterly for Exercise and Sport* 54, 169-178., 1983.
- [6] M. A. Sherman, A. Seth, and S. L. Delp. Simbody: multibody dynamics for biomedical research. *Procedia IUTAM Symposium on Human Body Dynamics*, 2(0):241 – 261, 2011.
- [7] M. Sun, G. W. Reed, and J. O. Hill. Modification of a whole room indirect calorimeter for measurement of rapid changes in energy expenditure. *Journal of Applied Physiology*, 76(6), June 1994.
- [8] K. I. Zaman, A. White, S. Yli-Piipari, and T. W. Hnat. K-sense: Towards a kinematic approach for measuring human energy expenditure. In *Proceedings of the 11th European Conference on Wireless Sensor Network (EWSN)*, 2014.

Demo Abstract: Object Recognition in Visual Sensor Networks. A comparison of two different operative paradigms.

A. Canclini, L. Baroffio, M. Cesana, A. Redondi, M. Tagliasacchi

*Dipartimento di Bioingegneria, Elettronica e Informazione, Politecnico di Milano
P.za Leonardo da Vinci 32, Milano, Italy
lastname@elet.polimi.it*

Abstract—This demo features a multi-hop Visual Sensor Network, capable of recognizing objects using two different visual paradigms. In the compress-then-analyze (CTA) paradigm, JPEG compressed images are transmitted through the network from a camera node to a central controller, where the analysis takes place. Conversely, in the analyze-then-compress (ATC) paradigm, the camera node extracts and compresses local visual features from the acquired images and transmits them to the central controller, where they are used to perform object recognition. We show that, in a bandwidth constrained scenario, the latter paradigm allows to reach higher application frame-rates, still ensuring excellent recognition results.

I. INTRODUCTION

Visual Sensor Networks (VSNs), composed of several inexpensive wireless camera nodes, are used to acquire images or video from the environment, which may be further processed to perform a broad range of complex visual analysis tasks, such as object recognition, event detection, localization and tracking, etc. The acquired visual information is typically delivered to a remote destination through low-power wireless transmissions (e.g., IEEE 802.15.4), and multiple visual or generic sensor nodes may be involved in the process of routing this information to its final destination. Due to their flexibility and low-cost, VSNs have attracted the interest of researchers worldwide in the last few years and are expected to play a major role in the evolution of the Internet-of-Things (IoT) paradigm. We focus on a fundamental application of Visual Sensor Networks: the detection and recognition of objects. This is typically accomplished by extracting global or local distinctive features from the acquired images, and by matching these features against a database of pre-stored and labelled features to finally perform object recognition. Such visual task can be practically implemented in different ways in the VSN depending on *where* in the network the task of feature extraction is performed. Broadly speaking, two distinct paradigms are possible: compress-then-analyze (CTA) and analyze-then-compress (ATC) [1].

Traditionally, the CTA paradigm is adopted (Figure 1(a)). The acquired images are locally compressed at the camera sensor (e.g., using standard techniques like JPEG) and

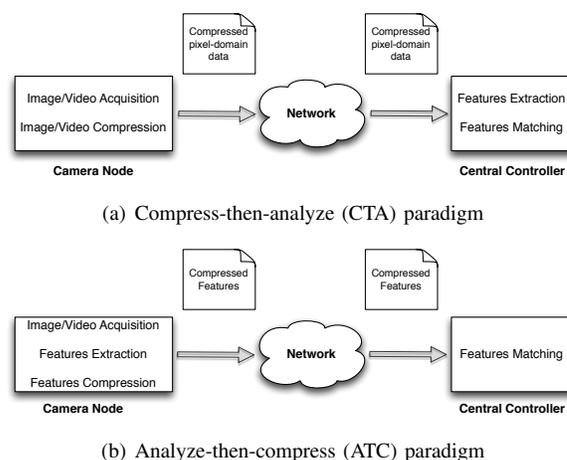


Fig. 1. The two different approaches to perform object recognition in visual sensor networks

delivered through the wireless sensor network to a central controller which performs object recognition. The bitstream flowing in the network thus includes the compressed version of a pixel-representation of the acquired image. Object recognition performed at the central controller is thus based on a compressed and lossy representation of the original image, which might significantly impair the recognition accuracy. Also, several works in the past demonstrated that multi-hop image transmission in VSNs results in high latency and low application frame rates, due to the struggling between bandwidth availability and requirements [2]. Alternatively, in ATC (Figure 1(b)) image features are extracted locally at the multimedia sensor, compressed, and then delivered to the final destination in order to enable higher level visual analysis tasks. The key tenet is that most visual analysis tasks can be carried out based on a succinct representation of the image, which entails both global and local features, while it disregards the underlying pixel-level representation. This approach allows to perform object recognition at higher frame rates with respect to the CTA case, as the visual information transmitted through the network is much more compact.

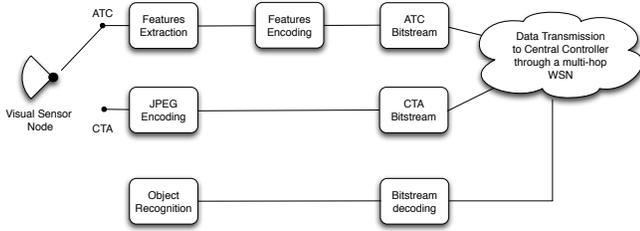


Fig. 2. Visual sensor node operational modes. Switching between ATC and CTA is remotely controlled by the user. The compressed multimedia data from the ATC or CTA paradigm is sent to a central controller where it is decoded and used to perform object recognition.

II. DESCRIPTION OF THE DEMONSTRATOR

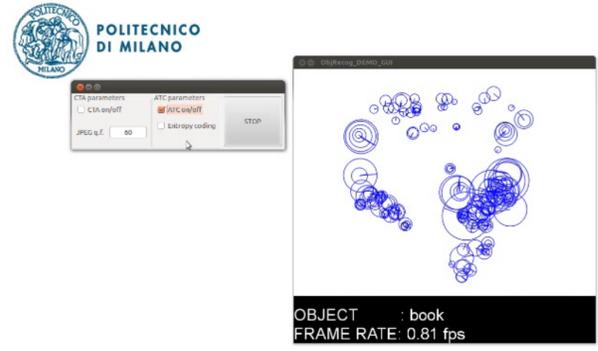
With reference to Figure 2, the demo testbed is composed of the following equipment:

- **Visual sensor node:** a battery-operated 720MHz ARM BeagleBone Linux computer which is geared with a Logitech USB camera to capture still images; the visual sensor node is also attached to a IEEE 802.15.4-compliant sensor node (TelosB platform or similar) to remotely transfer the visual content through low-power wireless links.
- **Network infrastructure:** a network of battery-operated IEEE 802.15.4-compliant TelosB sensor nodes which is used to route the visual information to a central controller.
- **Central controller:** a laptop with IEEE 802.15.4 communication capabilities to receive the multimedia content transferred by the visual sensor node and to perform visual analysis.

The visual sensor node is able to run both visual paradigms. Namely, as far as the CTA paradigm is concerned, the visual sensor node implements JPEG compression on the acquired image at different quality factors, which allows to trade off bandwidth with image quality. In the ATC case, the Binary Robust Invariant Scalable Keypoints (BRISK [3]) algorithm is implemented and used to extract local visual features. The camera node has also the ability to compress BRISK local features, by using a suitable encoding scheme to minimize the amount of bits required by each local feature without affecting the recognition capabilities [4]. At the central controller, the received features, in the ATC case, or the features extracted from the received JPEG image, in the CTA case, are matched against a database of labeled features, so that object recognition can be performed. The central controller also implements a graphical user interface which allows to visualize the result of object recognition and remotely control the operation of the visual sensor node, as illustrated in Figure 3; in details, the management console allows to change on the fly the type of paradigm, as well as all the most critical parameters of the two paradigms (quality factor for CTA, the number of features and the encoding scheme for the ATC). Moreover, in the ATC case, we implemented an engine to reconstruct a pixel-domain image from the knowledge of the local features only [5].



(a) Compress-then-analyze (CTA) paradigm



(b) Analyze-then-compress (ATC) paradigm

Fig. 3. Graphical user interface of the demonstration. (a) CTA paradigm: a JPEG compressed image is transmitted to the controller for recognition. (b) ATC paradigm: only a set of local visual features are remotely transmitted, ensuring recognition at higher frame rates.

The demo experiment will showcase first a classical object recognition applications with the testbed being able to recognize the type of object which is seen by the visual sensor node. Then, the testbed will be also used to assess the maximum frame rate, that is, the maximum number of images which can be processed per unit time, under the two paradigms, showing that ATC outperforms CTA.

III. DEMO VIDEO

A detailed video of the demonstration is available at www.greeneyesproject.eu.

REFERENCES

- [1] A. Redondi, L. Baroffio, M. Cesana, and M. Tagliasacchi, "Compress-then-analyze vs. analyse-then-compress: Two paradigms for image analysis in visual sensor networks," in *IEEE International Workshop on Multimedia Signal Processing*, 2013.
- [2] S. Paniga, L. Borsani, A. Redondi, M. Tagliasacchi, and M. Cesana, "Experimental evaluation of a video streaming system for wireless multimedia sensor networks," in *Med-Hoc-Net*, 2011, pp. 165–170.
- [3] S. Leutenegger, M. Chli, and R. Siegwart, "Brisk: Binary robust invariant scalable keypoints," in *Computer Vision (ICCV), 2011 IEEE International Conference on*, 2011, pp. 2548–2555.
- [4] A. Redondi, L. Baroffio, M. C. J. Ascenso and, and M. Tagliasacchi, "Rate-accuracy optimization of binary descriptors," in *IEEE International Conference on Image Processing 2013*, 2013, pp. 900–903.
- [5] E. d'Angelo, A. Alahi, and P. Vanderghyest, "Beyond bits: Reconstructing images from local binary descriptors," in *ICPR*, 2012, pp. 935–938.

Poster abstract: Sensor Network Based Surveillance System for Forest Protection

Gergely Vakulya, Gyula Simon, László Czúni, Péter Z. Varga
University of Pannonia, Veszprém, Hungary,
e-mail: {vakulya, simon}@dcs.uni-pannon.hu

Abstract—A sensor network based surveillance system is presented, which is able to detect illegal activities, especially timber theft, in forests. The sensor network is able to detect illegal entries of vehicles and also can detect ongoing logging activities. The sensors used in the system include accelerometers, acoustic, and magnetic sensors.

I. INTRODUCTION

ILLEGAL activities, especially the theft of timber cause damages in the order of ten million Euros each year for the forestry sector in Hungary. Although most forests are open and free to visitors and tourists, entry with vehicles (trucks, quads or motorbikes) is prohibited. Despite of the ban, illegal vehicular traffic, especially that of motorbikes and quads causes serious erosion and long term damages in the forest soil. Timber theft also causes major losses for the forestries. However, the detection of illegal activities and timely reaction is difficult: forest guards are overloaded and the limited number of personnel cannot efficiently guard large and remote territories. The proposed system helps the guards to detect illegal entries and logging activities, provides instant alarms with location information of the intrusion, to allow timely and efficient countermeasures.

Our proposed surveillance system uses inexpensive sensor networks to detect illegal entries of vehicles (either trucks to transport stolen timber or vehicles for dirt biking). Magnetic sensors and accelerometers are deployed along roads in the forest, mainly around the possible entry points, detecting any vehicular traffic. The logging activities are also detected, using acoustic sensors, which can detect the peculiar noise originating from chainsaws.

II. SYSTEM ARCHITECTURE

The block diagram of the system is shown in Fig. 1. Sensors are grouped into sensor islands, deployed in an area with diameter of typically 20-50 meters. The sensors in an island are connected to a central processing station through a GPRS capable base station. The central processing station, using the sensor readings and various data bases, provides status information and alarms for users.

This research was supported by the Hungarian Government and the European Union and co-financed by the European Social Fund under projects GOP-1.1.1-11-2011-0070 and TÁMOP-4.2.2.C-11/1/KONV-2012-0004. Gergely Vakulya was supported by the European Union and co-financed by the European Social Fund in the framework of TÁMOP 4.2.4. A/2-11-1-2012-0001 'National Excellence Program'.

In each sensor island various sensors are deployed: magnetic, accelerometer-based, and acoustic sensors are utilized, as shown in Fig. 2. The sensors are low power devices, with limited battery supply, while the base station has significant battery power and solar cells to recharge its battery.

The accelerometers and magnetic sensors are hidden beneath the surface, near roads, while the base and the acoustic sensors are deployed on trees.

III. SENSING

All sensors nodes are equipped with an ATmega128RFA1 combined microcontroller-radio chip, which is used for data processing and communication purposes. The sensor nodes are programmed in nesC, using the TinyOS operating system.

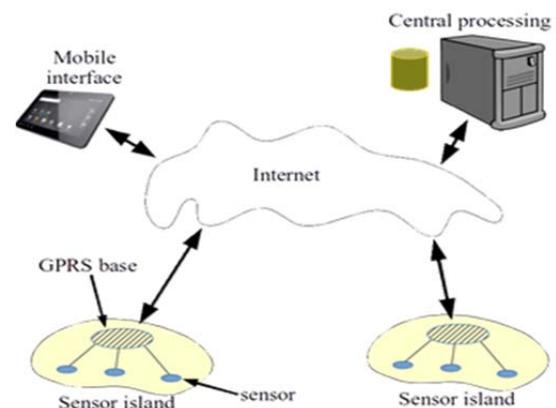


Fig. 1. The architecture of the surveillance system. Sensor islands are connected to the central processing station, where sensor fusion and event generation is performed.

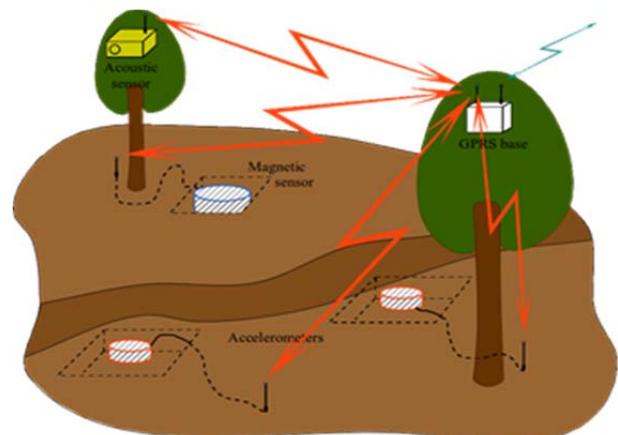


Fig. 2. Schematics of a sensor island, containing accelerometers, magnetic and acoustic sensors, and a GPRS enabled base station.

A. Magnetic detection

The body and engine of vehicles contain significant amount of ferromagnetic material, which disturbs the Earth's magnetic field, thus magnetic sensors can be used for vehicle detection (see e.g. [1]). The utilized sensors are MAG3110 three-axis digital magnetometers.

Due to the relatively short detection range of magnetic sensors, normally magnetic vehicle sensors are deployed under the roadway, below the passing-by vehicles. To prevent sensor damage in dirt-roads, sensors should not be placed under the road, but rather near to it. Our sensors are able to safely detect small (passenger) vehicles from 1.5 m, and large trucks from 2-3 m. Thus the sensors are hidden near the road, underground. To allow better concealment, the antennas are placed in concealed positions (e.g. nearby bushes) with underground cabling to the sensors (see Fig. 2).

B. Accelerometer-based detection

Passing-by vehicles cause significant amount of seismic vibration, which can be used for vehicle detection. In the system the BMA 180 three-axis accelerometer was used [2], the performance of which was enhanced by a spring and tiny mass, on which the accelerometer was fixed [4]. The detection range of the accelerometer-based detector is 4-8 m, depending on the actual road conditions and the size of the vehicle. The detector is placed underground in a metal container, few meters from the road, with concealed pigtail antenna similar to that of the magnetometer device.

C. Acoustic detection

Acoustic detectors use simple electret microphones with an appropriate amplifier and A/D converter to collect acoustic data. The acoustic detectors are used in two operation modes. In the chainsaw mode sensors are periodically switched on to collect 125 ms long sample records. In the autocorrelation function of the records signatures of chainsaw sound are searched for, using decision trees. When such signatures are found, an alarm message is sent to the base.

In the vehicle mode, acoustic sensors can be activated by the other sensors when vehicle detection occurred, to confirm the detection, using acoustic signatures. Our current acoustic sensors have an effective range of approximately 60 m.

D. Communication network

The sensor network inside a sensor island forms a star topology, the GPRS base station being always active, while sensors sleep as much as possible. To provide long lifetime for the sensors, communication overhead is reduced to minimum: sensors send short heartbeat messages in every 30 seconds, otherwise only alarms are sent. Sensors can also receive commands from the central processing station, using very small overhead by utilizing piggybacking on low level acknowledgement messages [3].

E. Sensor fusion and event detection

Sensor alarms, containing the type and time of detection, the sensor ID, and confidence values corresponding to the detection, are fused in the central processing station into actual

user alarms. The central station uses a data base on the individual sensor placement, describing the physical topology of the sensor network. When vehicles are detected, individual sensor alarms from a sensor island are compared to the expected alarm sequences, thus direction and speed can also be estimated, along with a confidence tag, reflecting how strongly the alarm is supported by sensor measurements.

The prototype monitoring screen is shown in Fig. 3, where the map of the monitored forest is shown along with the sensor islands. Actual alarms appear on the map and are also presented in table format in the lower part of the interface, showing the date and time of the alarm, the ID of the source sensor island, the type of the alarm (vehicle/chainsaw), related information (e.g. estimated velocity of the vehicle), and the associated confidence.



Fig. 3. Prototype monitoring screen of the surveillance system. Sensor islands are represented by colored bubbles. Detected events are listed below the map.

IV. CONCLUSION

A multisensory network was proposed to monitor forests to prevent illegal entries and timber theft. Accelerometer readings, magnetic, and acoustic data are fused to provide accurate alarm signals for the forest guard personnel. The system was designed to allow long lifetime in remote and harsh environment.

REFERENCES

- [1] B. Koszteczyk and G. Simon, "Magnetic-based vehicle detection with sensor networks", in *Proc. 2013 IEEE International Instrumentation and Measurement Technology Conference*, Minneapolis, USA, May 6-9, 2013, pp.265-270.
- [2] J. Smidla and G. Simon, "Accelerometer-Based Event Detector for Low-Power Applications," *Sensors*, vol. 13, no. 10, pp.13978-13997, Oct. 2013.
- [3] G. Vakulya and G. Simon, "Low-Power Communication Protocol for Low Duty Cycle Data Acquisition Applications", in *Proc. IEEE 2013 International Workshop on Measurements and Networking*, Naples, Italy, Oct. 7-8, 2013, pp. 58-62.
- [4] G. Vakulya, G. Simon, "Low Power Accelerometer Based Intrusion and Tamper Detector," In *Proc. 11th Int. Multi-Conference on Systems, Signals & Devices*, Castelldefels-Barcelona, Spain, Feb. 11 -14, 2014

Demo Abstract: Lightweight Continuous Indoor Tracking System

Zhuoling Xiao, Hongkai Wen, Andrew Markham, and Niki Trigoni

Department of Computer Science, University of Oxford. Email: firstname.lastname@cs.ox.ac.uk.

Abstract—Indoor tracking and navigation is a fundamental need for pervasive and context-aware smartphone applications. Although indoor maps are becoming increasingly available, there is no practical and reliable indoor map matching solution available at present. In this demo abstract, we describe a working prototype of a novel, robust and responsive tracking technique that is extremely computationally efficient (running in under 10 ms on an Android smartphone), does not require training in different sites, and tracks well even when presented with very noisy sensor data. The tracking system requires zero user effort (wandering, fingerprinting, etc.) – only the floor plan is required. We also demonstrate how it is able to accurately track the position of a user from accelerometer and magnetometer measurements only (i.e. gyro- and WiFi-free). We believe that such an energy-efficient approach will enable always-on continuous background localisation, enabling a new era of location-aware applications to be developed.

I. INTRODUCTION

Whereas GPS is the *de facto* solution for outdoor positioning, no clear solution has as yet emerged for indoor positioning despite intensive research and the commercial significance. Applications of indoor positioning include smart retail, navigation through large public spaces like transport hubs, and assisted living. The ultimate objective of an indoor positioning system is to provide continuous, reliable and accurate positioning on smartphone class devices [1]. We identify maps as the key to providing accurate indoor location. Based on a time-series of observations, such as inertial trajectories or RF scans, the goal is to reconcile the observations with the constraints provided by the maps in order to estimate the most feasible trajectory of the user, i.e. the sequence that violates the fewest constraints.

In this demo abstract, we describe a working prototype of continuous indoor tracking system that is lightweight and computationally efficient, but also robust to noisy data, allowing it to provide always-on and real-time location information to mobile device users. The proposed system uses an undirected graphical model, known as linear chain conditional random fields (CRFs) [2] which is particularly flexible and expressive, which allows us to capture correlations among observations over time, and to express the extent to which observations support not only states, but also state transitions.

The working prototype is very computationally efficient, running in < 10 msec on an Android phone, enabling real-time location computation online. It also offers high location accuracy even when it uses ultra low power sensors (e.g. accelerometer and magnetometer). With the trend of low power digital motion processors (DMPs), e.g. InvenSense MPU-6000/MPU-6050 which are able to task and process

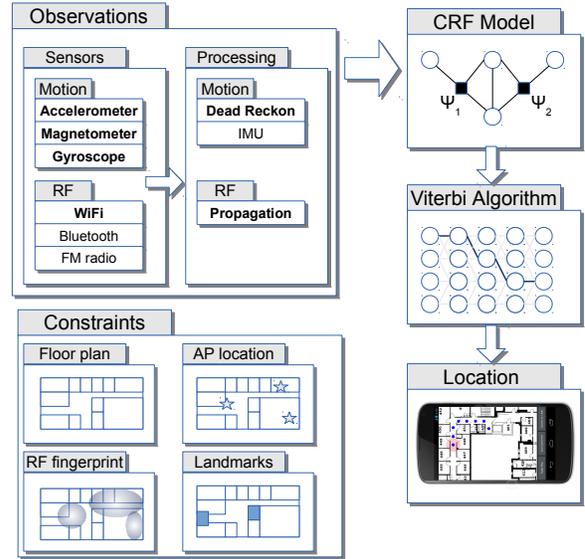


Fig. 1. Overview of the tracking system.

inertial data in bursts, while the system processor remains in a low-power sleep mode, our system could offer an always-on tracking service for a large range of low power devices.

II. TRACKING SYSTEM MODEL

The system architecture is shown graphically in Fig. 1, and is described through the use of an example.

When a user enters a building and launches the tracking application, the application requests a floor plan (along with other meta-data as generated by other systems, which could include fingerprint maps) from the server, if not already within the cache. Note that this is the only time that a user needs to reveal any data about their coarse position to a third party. The floor plan provides constraints over the set of possible positions a user can take, as well as allowed transitions between locations (i.e. a user cannot directly travel from one end of the building to the other without visiting intermediate locations). As such, the floor plan forms a sparse graph and can thus be efficiently stored in memory. Sensors on the user's phone collect data about the motion and (radio) environment. Motion sensors can include accelerometers, magnetometers and gyroscopes. Radio sensors can include WiFi, Bluetooth (low energy), FM radio and so forth. Raw sensor data is typically not immediately usable and needs to be processed. In the case of motion data, this could include dead reckoning trajectories based on counting steps and estimating heading, or

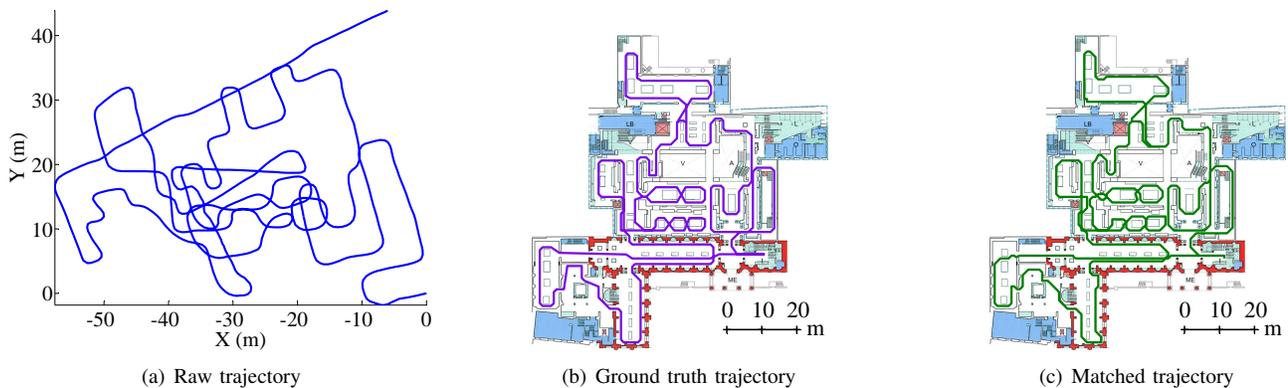


Fig. 2. Experiments in the museum (109m \times 89m), showing raw, ground-truth and matched trajectories.

using full IMU tracking in the case of foot mounted sensors. For RF data, a channel/propagation model can be used to relate received signal strengths to physical distances. Alternatively, raw signal strengths may be directly forwarded to the CRF model, to be later combined with RF fingerprint map data if available.

Maps and observations are combined using conditional random fields, an undirected graphical model which is particularly well suited to this sequential problem because it allows us to flexibly define *feature functions* that capture the extent to which observations support states and state transitions, given map constraints. As a user moves through the building, certain paths become unlikely, as they violate map constraints.

More specifically, this process involves four distinct steps:

Map pre-processing

This step takes a floor plan as input, and produces a graph that a) encodes a set of discrete states (locations), and b) represents physical constraints between discrete states imposed by the map. This information will then be fed to the second step, to help us define the CRF's states and feature functions.

Definition of states and feature functions

Our system uses feature functions to elegantly model different sensor data including the inertial measurements, RF measurements, visual measurements, user inputs, etc. and uses potential function to fuse them together;

Training to determine feature weights (optional)

With one or more true trajectories paired with respective sequences of sensor observations, training the CRFs to estimate weights is then performed by maximising the conditional probability of states given observations. Our system works well without training but training can help the tracking system to capture the special per-site features, which could help improve the tracking accuracy.

Inference to estimate location over time

The Viterbi algorithm is used to efficiently find the most likely sequence of states through the transition graph, culminating in an estimate of the user's location and quality thereof.

The first three steps are performed once for each building. The fourth step is performed online on the user's smartphone to track themselves.

Fig. 2 shows an example of our experimental results in

the museum, including the raw trajectory, ground truth, and matched trajectory, demonstrating an RMS error of 1.14m.

III. DEMO

In the demo session, we will show a working prototype of the continuous tracking system that offers accurate pedestrian location information without delay. The tracking system requires only the floor plan of the test site. In addition, to demonstrate the robustness of our system in different environments, we will also show the videos of our experiments in three other experimental sites: an office building, a museum, and a market.

IV. CONCLUSION

We demonstrated the merit of a novel continuous indoor tracking system, based on the application of conditional random fields. We have shown how it is robust, being able to operate with very noisy sensor data; lightweight, running in under 10 ms on a smartphone; and accurate, achieving the lowest RMS errors compared with other state-of-the-art approaches. Our system is able to establish a user's position using only dead-reckoned trajectories and a floorplan, without any external information such as a starting location or knowledge of WiFi access point locations. We believe that our tracking system has widespread application to a number of domains, as this single approach can be used with a wide variety of sensors and map information. One particularly relevant area is estimating location online and in real-time in resource-constrained body-worn sensors. In summary, we have presented a system that addresses the very pressing problem of providing accurate, low power, indoor tracking, that is responsive, robust and scalable.

Acknowledgements The authors would like to acknowledge the support of the EPSRC through grants EP/J012017/1 and the anonymous reviewers for their helpful comments.

REFERENCES

- [1] H. Wen, Z. Xiao, N. Trigoni, and P. Blunsom, "On assessing the accuracy of positioning systems in indoor environments," in *EWSN*, 2013.
- [2] J. D. Lafferty, A. McCallum, and F. C. N. Pereira, "Conditional random fields: Probabilistic models for segmenting and labeling sequence data," in *ICML*, 2001.

Demo Abstract:

Multi-Sensor Fusion for Indoor Pedestrian Tracking

Jan Medvesek, Andrew Symington and Stephen Hailes
 Department of Computer Science, University College London
 {j.medvesek,a.symington,s.hailes}@cs.ucl.ac.uk

I. INTRODUCTION

The outdoor pedestrian localization problem has in most cases been solved by the Global Positioning System (GPS). Positioning pedestrians indoors with a comparable accuracy remains an open problem. Solving this problem enables many new applications, including navigational aids, personnel tracking, location-aware mobile services and targeted advertising.

Consider a small sensor that is worn by a pedestrian on the waist, and used to perform localization. Resource constraints preclude the use of antenna arrays, cameras or laser scanners. However, small micro electro-mechanical (MEMS) sensors may be used to measure local motion, while and radio modules may be used to estimate range to fixed anchor points.

The objective of this demonstration is to show experimentally how motion and range measurements can be fused together to localize mobile sensors in cluttered environments.

II. BACKGROUND AND RELATED WORK

Pedestrian dead reckoning (PDR) describes a method by which inertial measurements are filtered to obtain an estimate of the path taken by a person walking. Approaches can be divided into full inertial navigation systems (INS) or step and heading systems (SHS) [1]. The INS approach typically uses a full six degree-of-freedom kinematic filter to maintain an estimate of 3D position, velocity and orientation. Acceleration is integrated twice through orientation to update position, and so the resulting position error grows quadratically with time. As a result of the large noise exhibited by MEMS sensors, even well-calibrated systems can yield position drift of ten meters or more within the first few seconds of operation. To mitigate this problem systems typically identify step events by thresholding acceleration or angular velocity [2], and exploit the fact that the foot is stationary during this interval to reset velocity to zero [3]. The SHS approach estimates stride length¹ and uses an attitude and heading reference system (AHRS) to maintain an estimate of heading; the geogravitational and geomagnetic vectors provide a reference frame against which orientation can be corrected. Step events are detected and stride length is integrated through heading to update position.

Radio measurements can be used to periodically correct the drift-corrupted position. The relationship between received signal strength (RSS) and pairwise range in free-space may be described by a log-distance path loss model [4]. Pairwise

¹Stride length differs from person-to-person, but can be approximated for an individual from step frequency, gender and height.

range is also related to time-of-flight (TOF), which in the general case is difficult to measure, because of clock frequency instability. Nevertheless, chirp spread spectrum can be used in conjunction with a symmetric double-sided two-way ranging algorithm to provide centimeter-resolution ranging in free space [5]. Depending on the network topology, one of a number of different parametric approaches can be used to convert multiple RSS or TOF range estimates to a position estimate. In practice however the accuracy is limited, as environmental clutter causes signal reflection, absorption and multipath. State-of-the-art fingerprinting algorithms exploit this spatially-correlated signal variation to model the posterior distribution of RSS over space, using training data [6].

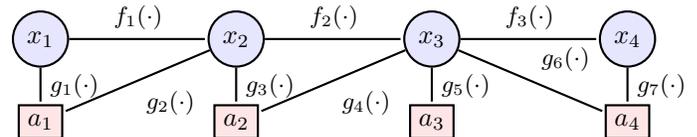


Figure 1. The localization problem expressed as a graphical model. Circles represent unknown sensor positions, while squares represent known sensor positions. Measurements are shown as edges that constrain relative states.

III. PROPOSED LOCALIZATION ALGORITHM

The localization problem can be represented by the connected graph shown in Fig. 1. The objective of localization is to estimate values for all *unknown* sensor positions $x_1 \dots, x_n$, given *known* anchor positions a_1, \dots, a_m , *processes* $f(\cdot)$ that relate two sequential sensor positions, and *corrections* $g(\cdot)$ that relate a sensor's position to a subset of anchor through TOF, RSS or magnetic field strength measurements.

A. Recursive Solution

The proposed recursive solution fuses motion and radio measurements using a cascaded filter. The upper part of the filter implements an SHS for estimating displacements from inertial measurements. On every step event the upper part passes a displacement measurement to the lower part. The lower part of the filter uses a Sequential Monte Carlo (SMC) approach for estimating position, by fusing the displacements with TOF, RSS and magnetic field strength measurements.

The SHS estimates displacements from accelerometer, gyroscope and magnetic measurements. A gradient descent algorithm [7] is used to estimate bearing. Steps are triggered no



Figure 2. An electronic theodolite is used to obtain a millimeter-resolution ground truth. Four anchors are placed on roof runners, and a single tag sensor is carried around the room. The sensors records inertial, geomagnetic, time-of-flight and received signal strength measurements for processing offline.

faster than 1Hz, and when net angular rate energy drops below 0.6 radians per second. At each step a displacement vector is calculated by integrating stride length through bearing.

The SMC approach uses a particle filter that maintains a multi-modal estimate of position. At each time step a particle's position is first updated with a displacement from the SHS. The estimate is then corrected using RSS, TOA and magnetic field strength measurements, when available. A Gaussian processes describes the posterior distribution of each quantity, and particles are resampled to maximize measurement likelihood.

B. Full solution

Assume that the set of all unknown sensor positions is denoted $X = \{x_1, \dots, x_n\}$. Processes $f(\cdot)$ and corrections $g(\cdot)$ are bundled together into a measurement set M . The measurement model $h_i(\cdot)$ for the i th measurement extracts the relevant states from X to *predict* the measurement value. A residual is then calculated using the observation z_i (a displacement or range estimate) and the measurement noise covariance R_i . Localization is carried out by assigning values to X in such a way that minimizes the sum of the weighted square residuals, through the objective function below.

$$\bar{X} = \arg \min_X \sum_{i \in M} (z_i - h_i(X)) R_i^{-1} (z_i - h_i(X))$$

It is often the case that $h_i(\cdot)$ functions are non-linear. Most solvers therefore typically carry out a first-order linearization to obtain a sparse Jacobian that relates state values. The block structure of this Jacobian can be exploited to improve solver performance [8]. Linearization is an approximation and therefore such solvers are non-optimal, which means that they may converge to local minima. Proper initialization is key to obtaining an accurate solution, and for this reason the proposed method uses the recursive solution as an initial state estimate.

IV. EXPERIMENTAL STUDY

The proposed algorithm will be demonstrated experimentally using one sensor and four anchors. Each device is based around the *Orisen Prime* development board. This board

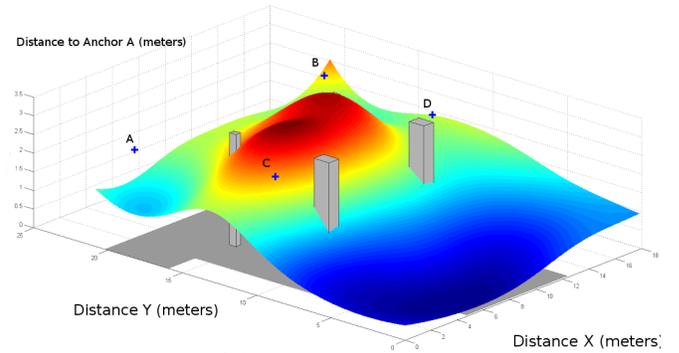


Figure 3. A Gaussian process describing the measured TOF range from anchor A. The presence of pillars in the office positively biases the measurement.

features a microcontroller, flash storage, an IEEE 802.15.4 radio, inertial measurement unit and other sensors. The board was coupled with a *nanoPAN* radio to provide TOF estimates.

Fig. 2 shows a sensor being carried by a pedestrian within an office space, while being tracked by an electronic theodolite to obtain a ground truth. In the first run data was collected to fit Gaussian process parameters. The resulting measurement models were then used to perform localization for a second experiment run. The Gaussian process describing the *measured* TOF range from all points in the X-Y plane from anchor A is shown in Fig. 3. Note that the presence of a pillar in the room prevents line-of-sight, positively-biasing the TOF range. More results will be presented in the demonstration.

ACKNOWLEDGMENTS

The work in this proposal was partially funded by the European Commission's Seventh Framework Programme for Cognitive Systems & Robotics. The authors would like to extend their gratitude to Graeme McPhilips for the design of the *Orisen Prime* board, to Dietmar Backes for assisting with capturing ground truth measurements using an electronic theodolite, and to Simon Julier for guidance with estimation.

REFERENCES

- [1] R. Harle, "A Survey of Indoor Inertial Positioning Systems for Pedestrians," *IEEE Comm. Surveys & Tutorials*, vol. 15, no. 3, pp. 1–13, 2013.
- [2] I. Skog, P. Händel, J. Nilsson, and J. Rantakokko, "Zero-velocity detection - an algorithm evaluation.," *IEEE transactions on biomedical engineering*, vol. 57, pp. 2657–2666, July 2010.
- [3] E. Foxlin, "Pedestrian tracking with shoe-mounted inertial sensors," *IEEE Computer Graphics and Applications*, vol. 25, no. 6, pp. 38–46, 2005.
- [4] N. Patwari, J. Ash, S. Kyperountas, A. Hero, R. Moses, and N. Correal, "Locating the nodes: cooperative localization in wireless sensor networks," *Signal Processing Magazine*, vol. 22, pp. 54–69, July 2005.
- [5] C. Rohrig and M. Muller, "Localization of sensor nodes in a wireless sensor network using the nanoLOC TRX transceiver," in *IEEE 69th Vehicular Technology Conference.*, pp. 1–5, 2009.
- [6] B. Ferris, D. Hähnel, and D. Fox, "Gaussian processes for signal strength-based location estimation," in *Robotics Science and Systems*, (Philadelphia, PN), 2006.
- [7] S. Madgwick, A. Harrison, and A. Vaidyanathan, "Estimation of IMU and MARG orientation using a gradient descent algorithm.," in *The International Conference on Rehabilitation Robotics*, vol. 2011, (Zurich, Switzerland), p. 5975346, Jan. 2011.
- [8] M. Lourakis and A. Argyros, "SBA: a software package for generic sparse bundle adjustment," *ACM Transactions on Mathematical Software*, vol. 36, pp. 1–30, Mar. 2009.

Poster Abstract: A Case for Magneto-Inductive Indoor Localization

Traian E. Abrudan, Andrew Markham, and Niki Trigoni

Department of Computer Science, University of Oxford. E-mails: firstname.lastname@cs.ox.ac.uk

Abstract—In this work, we show that the magneto-inductive (MI) indoor localization using tri-axial coils possesses unique advantages over the RF-based techniques. The spatial distribution of a generated magnetic field strength can be easily predicted using simple linear models. In addition, the field strength is very stable over time, even in the presence of moving people, which makes MI localization a promising indoor positioning technology.

I. INTRODUCTION

In the past years, there has been a growing interest in localization systems using magnetic fields [1]–[7]. This is partly because cheap magnetic sensors are nowadays available in almost every hand-held smart device. In addition, these systems usually operate at extremely low frequencies, in the near field region. Consequently, the corresponding signals are non-propagating, and they do not experience multipath and shadow fading, such as the high frequency radio waves. Unlike microwaves, the magnetic fields experience no absorption by water (e.g. by the human body), which makes them very attractive for localization inside highly populated buildings, where people are continuously moving [8]. Moreover, they do not usually require line-of-sight between devices, as they have the ability to penetrate through soil, concrete and rock with negligible attenuation [2]. Existing magnetic localization approaches use either indigenous magnetic fields, such as the Earth’s magnetic field and/or the magnetic fields generated by home electronics [5]–[7], or low-frequency magnetic fields generated locally, in purpose for localization [1]–[4]. The major challenge in indoors is accounting for the distortions in the magnetic field, in the vicinity of massive metallic objects such as building structures, elevators, etc. In spite of that, we provide few results showing that the magnetic field strength is much more predictable than the strength of the radio waves.

II. SYSTEM DESCRIPTION

Our measurement system consists of one magnetic transmitter (TX) and one magnetic receiver (RX) that operate at a carrier frequency is 2.5KHz, and at symbol rate of 32 symbols/s. The Received Signal Strength Indicator (RSSI) is estimated by measuring the energy at the receiver, corresponding to a known preamble of 120 symbols. One particularity of our system is that both TX and RX are equipped with tri-axial coils, which makes the RSSI invariant to the relative rotations of TX and RX [1], [4]. RX is connected to a computer via the USB port, where the data is stored and processed offline.

III. EXPERIMENTAL RESULTS

We performed a fingerprinting experiment in a large lecture room. Fig. 1 (left column) shows the measurement scenario, with TX at the origin (black square marker) and RX moving on the horizontal X - Y grid (black dot markers), whose step

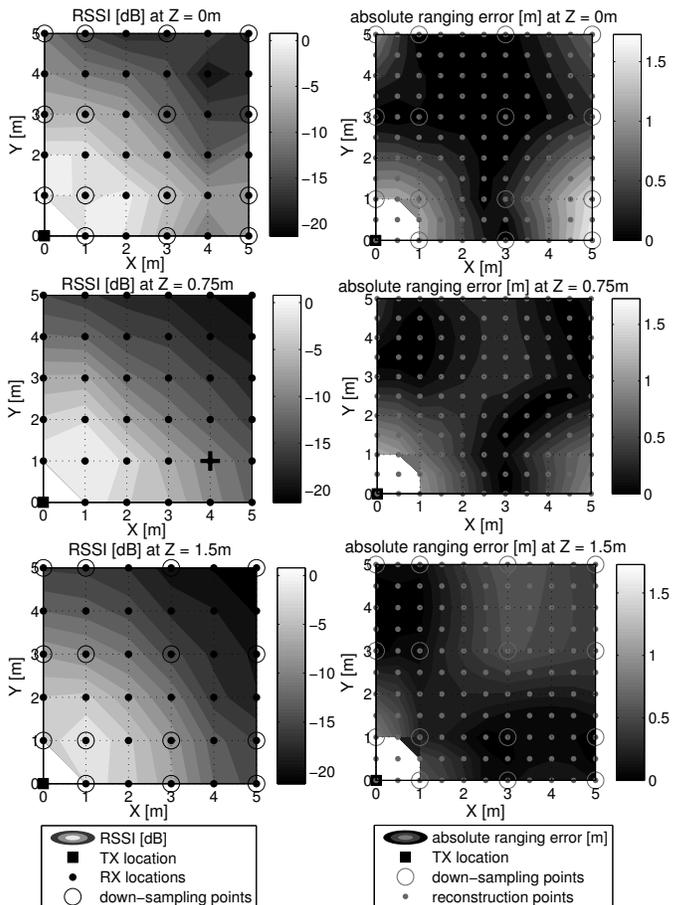


Fig. 1. Sparse magnetic fingerprinting (RSSI – left column) and ranging using the reconstructed RSSI (error in the range estimate $\hat{r} - r$ – right column).

is 1 meter in each dimension. Three different heights were considered: $Z \in \{0, 0.75, 1.5\}$ meters, respectively. The three plots in the left column of Fig. 1 show the RSSI magnetic map for each horizontal “slice”. We may notice that the constant RSSI curves are approximately circular (with small variations at the floor level). The signal strength is easily predictable, unlike for example the WiFi signals that are subject to severe fading. Therefore, the range estimation can be done using simple models. In Fig. 2, we show how the overall RSSI decreases with distance. The red dots represent the measured RSSI values at all points of the rectangular lattice used in the fingerprinting experiment. Few extra-measurements (indicated by black “x” marker) were taken at larger distances in order to show the RSSI trend. The model is estimated from the measurements using linear Least-Squares (LS), and is shown in Fig. 2 by the continuous thick gray line. We may notice

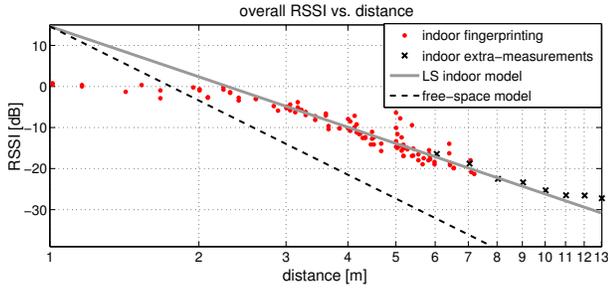


Fig. 2. The decay of RSSI with distance r . Measurements vs. models.

that the slope corresponding to the derived indoor model is about 40dB/decade, compared to the free-space model whose slope is 60dB/decade [2], shown in Fig. 2 by the black dashed line. Therefore, the magnetic field strength decays slower in indoors. This phenomenon might be caused by the fact that the surrounding ferrous materials act as passive re-radiators [1].

In order to prove the spatial predictability of the RSSI, we down-sampled the 3 slices approximately by a factor of 2 along X and Y axes. In the vertical dimension, the middle horizontal slice was eliminated completely. The down-sampling points are marked in Fig. 1 with black/gray circles. We reconstructed the missing data using simple 3D linear interpolation of the down-sampled data with a resolution of 0.5 m in all three dimensions. The reconstruction points and the corresponding ranging errors are shown in the right column of Fig. 1 (only 3 out of 7 interpolated horizontal slices are shown). Despite the sparse down-sampling, the signal is predicted with an accuracy of few dB. Slightly larger ranging errors occur at the floor level, but overall, the signal is much easier to predict than, for example, WiFi. The overall ranging bias using the interpolated RSSIs and our LS model is $b_r = E\{\hat{r} - r\} = 0.09$ meters, whereas the standard deviation of the ranging errors is $\sigma_r = [E\{(\hat{r} - b_r)^2\}]^{1/2} = 0.45$ m. The maximum ranging error is $e_{\max} = 1.73$ m at $Z = 0$, and only 0.88m at $Z = 1.5$ m. This shows that ranging models for magnetic localization are very reliable and that compared to the radio maps, the magnetic maps are much easier to reconstruct from spatially sparse samples. Similar results were obtained when the range was estimated using the LS model in Fig. 2 only (no fingerprinting) and the RSSIs: $b_r = 0.14$ m, $\sigma_r = 0.53$ m, and $e_{\max} = 1.73$ m. The slightly larger variance might be caused by the fact that the range estimation was done in a per-measurement basis, the spatial correlation being neglected. The RSSI is easily predictable, and this is particularly important when fingerprinting-based localization is used. Much of the tedious map construction work can be avoided by using sparse sampling.

Another important advantage of MI localization is that the generated magnetic field strength is very stable over a long period of time, unlike RF-based techniques. In Fig. 3, we show the variation of the RSSI of the magnetic link versus the RSSI corresponding to a WiFi link. RX and TX were stationary, placed at the same location in both cases, about 4 meters apart, with TX at the origin. The location of RX from which the RSSI time variation was analyzed is marked by a black cross in middle subplot of the first column in Fig. 1 (corresponding to $Z = 0.75$ meters). The RSSI was recorded for a period of approximately 8 minutes for each of the transceivers, while two persons were walking between TX and RX, crossing the link, and approaching TX and RX. We may notice that the RSSI corresponding to the magnetic transceiver has negligible

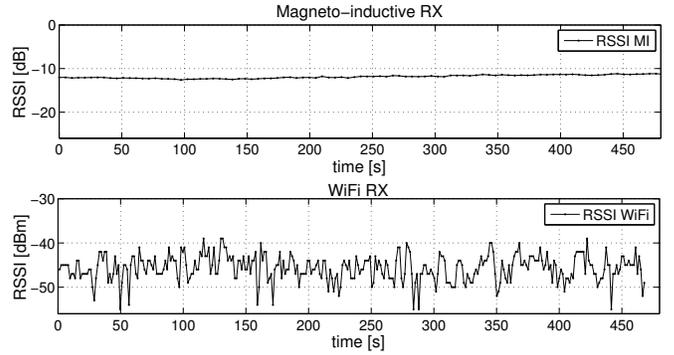


Fig. 3. The RSSI stability over time. Magnetic link vs. WiFi link.

variation compared to the RSSI of the WiFi. A similar analysis is given in [8], but for the Earth's field.

IV. CONCLUSION

In this paper, we show that when tri-axial coils are used both at TX and RX, the indoor magnetic RSSI is very easy to predict both in space and time, unlike the RF RSSI used in most of the wireless standards. The fast decay of the magnetic field is undesirable from the point of view of the transmission range, but it allows to distinguish between very closely-spaced distances. This has been shown in [2] for underground animal tracking. Our future work will focus on RSSI fingerprinting independently in each axis. We aim to achieve full 3D location and orientation estimation using a single transmitter.

V. ACKNOWLEDGEMENTS

The authors would like to thank EPSRC for funding this research (Grant ref. EP/L00416X/1 Digital Personhood: Being There: Humans and Robots in Public Spaces (HARPS)).

REFERENCES

- [1] A. Markham and N. Trigoni, "Magneto-inductive networked rescue system (MINERS): taking sensor networks underground," in *11th International Conference on Information Processing in Sensor Networks (IPSN 2012)*, Beijing, China, Apr. 16–20 2012.
- [2] A. Markham, N. Trigoni, S. A. Ellwood, and D. W. Macdonald, "Revealing the hidden lives of underground animals using magneto-inductive tracking," in *8th ACM Conference on Embedded Networked Sensor Systems (Sensys 2010)*, Zürich, Switzerland, Nov. 2010.
- [3] A. Sheinker, B. Ginzburg, N. Salomonski, L. Frumkis, and B. Kaplan, "Localization in 3-D using beacons of low frequency magnetic field," *IEEE Transactions on Instrumentation and Measurement*, no. 99, 2013.
- [4] G. Pirkel and P. Lukowicz, "Robust, low cost indoor positioning using magnetic resonant coupling," in *ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp 2012)*, Pittsburgh, USA, 2012.
- [5] M. Frassl, M. Angermann, M. Lichtenstern, P. Robertson, B. J. Julian, Marek, and Doniec, "Magnetic maps of indoor environments for precise localization of legged and non-legged locomotion," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2013)*, Tokyo, Japan, Nov. 2013.
- [6] J. Haverinen and A. Kemppainen, "Global indoor self-localization based on the ambient magnetic field," *Robotics and Autonomous Systems*, vol. 57, pp. 1028–1035, 2009.
- [7] S. Rahok and O. Koichi, "Odometry correction with localization based on land-markless magnetic map for navigation system of indoor mobile robot," in *4th International Conference on Autonomous Robots and Agents (ICARA 2009)*, Feb. 2009, pp. 572–577.
- [8] B. Li, T. Gallagher, A. G. Dempster, and C. Rizos, "How feasible is the use of magnetic field alone for indoor positioning?" in *2012 International Conference on Indoor Positioning and Indoor Navigation (IPIN 2012)*, Guimarães, Portugal, 13–15 Nov. 2012.

Poster Abstract: WiFi Sensors Meet Visual Tracking For An Accurate Positioning System

Savvas Papaioannou, Hongkai Wen, Zhuoling Xiao, Andrew Markham, and Niki Trigoni
Department of Computer Science, University of Oxford. Email: firstname.lastname@cs.ox.ac.uk.

Abstract—In this poster abstract, we propose a new positioning technique that can localize people by combining WiFi information from their mobile devices with visual tracking. We show that the proposed approach can improve visual tracking by resolving motion and appearance ambiguities while at the same time can uniquely identify each person with their device ID.

I. INTRODUCTION

The development of practical and accurate indoor positioning systems has received a lot of attention from the research community. One commonly used positioning technique is vision-based tracking. Many existing vision-based tracking algorithms are capable of accurately estimating the trajectories of multiple targets from video footage. However, the estimated trajectories are usually *anonymous*. Visual identification (e.g. face recognition) cannot always be applied since it requires knowledge on the mapping between IDs and pictures, and can be computationally expensive and privacy intrusive. On the other hand, the WiFi-based positioning systems are able to localize the mobile devices carried by the targets. Therefore, the WiFi systems possess the knowledge of the targets' ID, but typically require a stable radio map, use many APs and often a floor plan to achieve good accuracy.

Motivated by the above problems, we propose a new positioning technique which is able to perform accurate localization and privacy-preserving identification. The key idea is to exploit the existing WiFi and camera infrastructure, which is available in most of today's large indoor environments. We use visual tracking techniques to detect moving objects in the camera footage, and generate anonymous tracklets based on a motion model. The tracklets are then fused with WiFi signal strength measurements to produce accurate trajectories of each target. Unlike most of the existing work which uses WiFi information to *match* the trajectories produced by visual tracking, our approach incorporates WiFi measurements in visual tracking to *generate* the trajectories, and thus can achieve similar performance with less infrastructure and noisy signals.

II. PROPOSED APPROACH

A. System Architecture

The proposed approach contains two components: a *foreground detector* and a *tracker*, as shown in Figure 1. We assume the indoor environment is covered by one calibrated stationary camera, and multiple WiFi access points (APs) with known locations are deployed. A number of people (targets) are moving around with their mobile devices

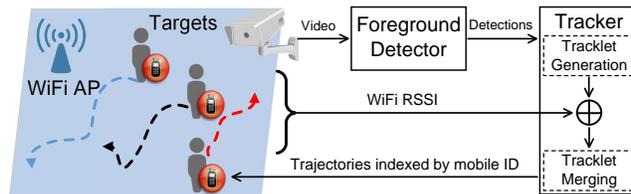


Fig. 1. The architecture of the proposed system.



Fig. 2. The detections generated by our implementation of foreground detector.

receiving WiFi beacons from the APs periodically. The captured video stream is processed by the foreground detector to extract detections of moving objects, and then the tracker fuses the visual detections with WiFi RSSI measurements collected by the mobile devices to a) generate accurate trajectories, and b) identify the target ID of each trajectory.

B. Foreground Detector

The foreground detector detects any moving objects in the camera footage, and can be implemented in a number of different ways [1]. Comparing to the more sophisticated object detection techniques (e.g. [2]), foreground detector is lightweight, and it can be used in real-time in embedded camera networks [3]. However, the detections generated by a foreground detector can be very *noisy*. We have observed three distinct cases in our preliminary experiments where this happens (as shown in Figure 2): a) multiple detections are generated for one moving object (D1 and D2), b) a detection contains no moving object at all (D4), and c) one detection contains multiple moving objects (D5).

C. WiFi-based Tracker

The key component of the proposed approach is the WiFi-based tracker, which works in two steps: 1) *tracklet generation* and 2) *tracklet merging*. Concretely, suppose we have m targets. The input video is divided into s segments, each of which has n frames. The tracker first generates short trajectories (tracklets) for the s video segments, and

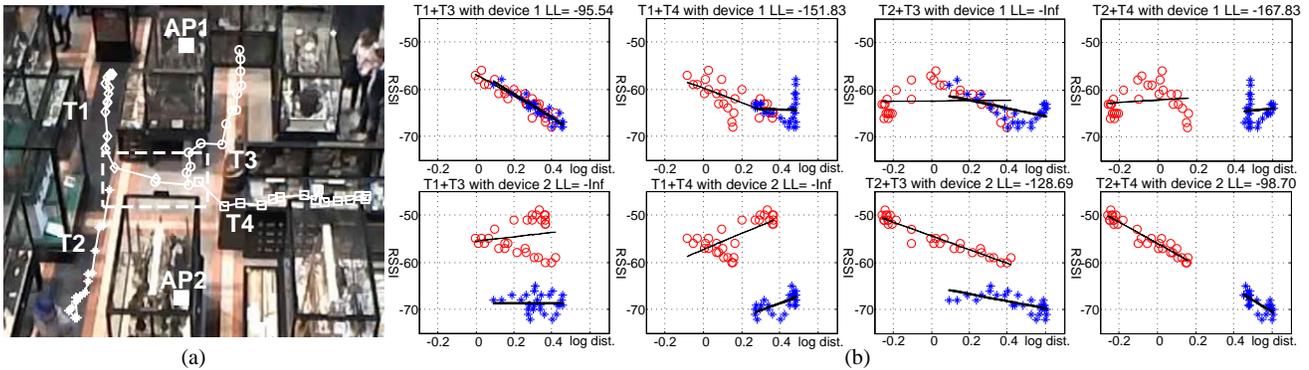


Fig. 3. (a) Tracklets generated for two targets in two video segments, where $T1, T2$ are in the first segment, and $T3, T4$ in the second. (b) WiFi signal strength measurements can help find the correct solutions: $T1+T3$ with device 1 and $T2+T4$ with device 2 (the first plot and the last plot).

then merges them to produce the final trajectories of the m targets.

Tracklet generation: For each segment a weighted directed acyclic graph $G = (V, E, W)$ is generated, where the vertices V represents the detections computed by the foreground detector, and are partitioned into n disjoint groups according to frames. The edges E are defined between the vertices in different groups to form a trellis diagram. The weight w of an edge e describes the appearance and motion costs that two detections belong to the same trajectory. The appearance cost is determined by the color histogram intersection between the two detections. The motion cost is calculated by comparing the distance between the two detections with a motion model learned from the data, which constrains the movement speed. For one target, we use a Viterbi-like technique to traverse the trellis graph to find the tracklet with the minimum cost by selecting one detection from each group of vertices. We then remove the found tracklet from the graph and perform the same procedure for the next target until all m tracklets are found.

Tracklet merging: In this step, the tracker merges the tracklets generated for each segment to produce the complete trajectories, and assigns the correct target IDs to them. We use a sliding window algorithm, which considers the tracklets in k segments at a time. Concretely, the algorithm performs k -partite matching to find the most likely trajectory of a target. It iteratively searches the space of all possible solutions that contain k tracklets within the current window, guided by a cost function: $C(l) = w_A C_A(l) + w_M C_M(l) + w_W C_W(l)$, where l is a solution (i.e. a possible trajectory associated with a target ID), C_A, C_M, C_W are the cost functions for appearance, motion and WiFi respectively, and w_A, w_M, w_W are the normalizing weights. $C_A(l)$ is computed from the pairwise intersections of the average color histograms of tracklets in l . $C_M(l)$ is evaluated by comparing the pairwise distances between the tracklets in l (begin and end points) with the motion model. The WiFi cost $C_W(l)$ describes how consistent is the solution l with the WiFi RSSI measurements of a device (carried by a target). Given the locations of the APs, the log-distances between the points on the solution l and the APs can be known exactly. If l is correct, then the relationship between the log-distances and the sequence of signal strength measurements should be *linear*, governed by

the radio propagation model. With this intuition, our tracker performs Bayesian linear regression on the log-distances and the RSSI measurements, and $C_W(l)$ is defined as the log-likelihood that the observed RSSI sequence agrees with the trajectory l under the radio model. Figure 3 shows how WiFi can help to associate the tracklets correctly. We consider two targets (two different devices) in two video segments, where $(T1, T2)$ and $(T3, T4)$ are the tracklets in the 1st and 2nd segment respectively. In this case it is very difficult to find the correct solution based only on motion (assuming similar appearance) since the targets are close in the highlighted region. Figure 3(b) shows the relationship between the log-distances and the observed RSSI measurements from 2 APs for all possible solutions (2 devices \times 4 possible trajectories), where only the two solutions: a) $T1+T3$ with device 1 and b) $T2+T4$ with device 2 are correct (also indicated by the higher log-likelihood (LL)). Note that the same technique can also be used to deal with occlusions (suppose no trajectory was in the highlighted box in Fig. 3(a)), since WiFi measurements can correctly connect the tracklets to fill the gap. Our initial results in a real setting (a museum) indicate that the proposed technique can be used to resolve motion and appearance ambiguities.

III. ACKNOWLEDGMENT

We would like to thank Laing O'Rourke for funding this research and also the Pitt Rivers Museum for allowing us to conduct our experiments in the museum's space.

IV. CONCLUSION

We propose a new positioning system that integrates WiFi information with visual tracking. The novelty of our approach is that it leverages WiFi measurements to improve the performance of visual tracking, and offer accurate localization and identification at the same time.

REFERENCES

- [1] C. Stauffer and W. E. L. Grimson, "Adaptive background mixture models for real-time tracking," in *Proc. CVPR*, 1999.
- [2] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan, "Object detection with discriminatively trained part-based models," *PAMI*, vol. 32, no. 9, pp. 1627–1645, 2010.
- [3] Y. Shen, W. Hu, J. Liu, M. Yang, B. Wei, and C. T. Chou, "Efficient background subtraction for real-time tracking in embedded camera networks," in *Proc. SenSys*, 2012.

Poster abstract: Robust Phase Estimation for RSSI signals

Gergely Zachár, Gyula Simon
University of Pannonia, Veszprém, Hungary
e-mail: {zachar, simon}@dcs.uni-pannon.hu

Abstract—A distributed noise-tolerant and accurate phase-difference measurement scheme is proposed, which can be run on low end devices with limited computational power. The proposed algorithm can be used as a key building block in radio interferometric localization systems. The operation of the algorithm is illustrated by real measurements.

I. INTRODUCTION

An emerging localization technique in sensor networks uses radio-interferometry [1]-[3]. In such applications precise distributed phase-difference measurements are necessary: two radio receivers detect interferometric radio signals (the envelope of which is the received signal strength, RSSI signal) at the same time (see Fig. 1). The localization and tracking algorithms use the difference between the phases of the two RSSI signals (see e.g. [1]). In this paper a robust and accurate method to measure such phase difference is proposed. The proposed method has low computational complexity on low-end devices and also provides fast and accurate measurements.

II. PHASE MEASUREMENT

In large scale sensor-networking applications, where a large number of nodes are utilized, the nodes are usually inexpensive devices with limited computational power. To perform distributed phase difference measurements on such devices, the phase estimation algorithm should fulfill the following requirements:

1. The processing requirements should be low enough to enable the utilization of the algorithm in devices with limited resources and processing power: the nodes used for measurement are based on an 8-bit microcontroller.
2. The processing should be fast, since the measurements should be done as frequently as possible.
3. The algorithm must handle the logarithmically-scaled sinusoidal signal form, since the RSSI values measured in the radio are provided on a logarithmic scale.
4. The algorithm must be robust to handle distorted and noisy signals.

This research was supported by the Hungarian Government and the European Union and co-financed by the European Social Fund under project TÁMOP-4.2.2.C-11/1/KONV-2012-0004. Gergely Zachár was supported by the European Union and co-financed by the European Social Fund in the framework of TÁMOP 4.2.4. A/2-11-1-2012-0001 'National Excellence Program'.

Based on the requirements a tailor-made algorithm was designed, which is using minimal resources on the devices.

To allow fast and still accurate processing, low-end devices use only integer calculus for intermediate results, while the final result is calculated on a host computer with floating-point arithmetic. Thus the measurement result can be more accurate while soft floating-point operations are avoided on the devices.

The operation of the proposed method can be seen in Fig. 2. First the RSSI samples are collected: the sampling frequency and the length of the record are determined so that at least a few full periods are present in the data record. Data recording is performed in a strictly synchronized manner: the two receiver nodes start the recording when a start message is received (using reference broadcast). Then in both data records reference points are searched for.

Note that the signal can be rather asymmetric, due to its logarithmic nature. Signal maximum or minimum points are good choices for reference, since their positions are invariant to the actual shape (see Fig. 1). The problem with these points is that the signal's slope is zero here, thus direct measurement is inaccurate in the presence of noise. Other reference points with higher slope (e.g. signal crossing a given value) would be easier to measure but they are not invariant to the actual signal's shape, resulting offset in the phase measurement. The proposed solution uses maximums as reference points but the estimation of these positions is done indirectly, resulting in an accurate and noise-resilient solution.

In the first processing step the average of the signal is calculated (red line). Then the upper and lower hysteresis

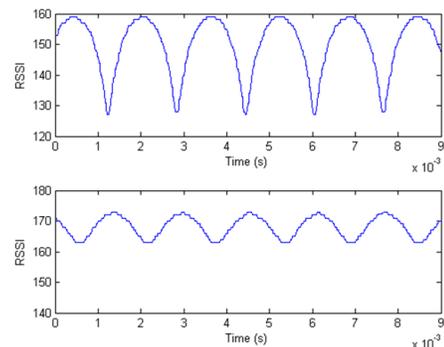


Fig. 1. Typical interferometric RSSI signal forms, registered at two receivers

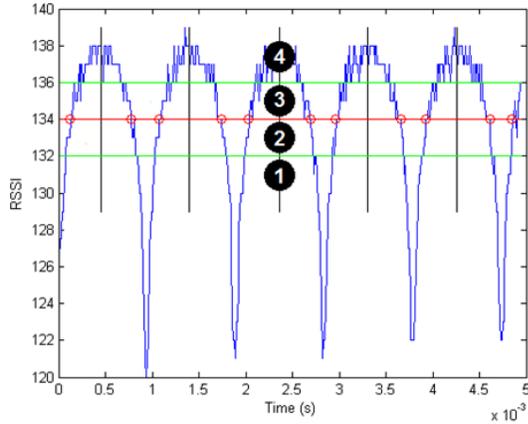


Fig. 2. Illustration of the phase measurement: red line: mean value, green lines: hysteresis limits, 1-4: amplitude ranges, red circles: 2-3 and 3-2 transitions, and black lines: estimated times of the reference points

limits (green lines) are determined, which have distances of Δ from the average, where value Δ is determined to be higher than the maximal noise level. The average and the two limits divide the amplitude space into four regions, as shown in Fig 2. The further processing of the signal is performed with a state-machine, which follows the signal through regions 1-2-3-4-3-2-1-..., in this strict order. When transitions 2-3 and 3-2 occur, the times of the transitions are marked (t_{23} and t_{32} , respectively, zero being at the beginning of the data record) and the time of the maximum of the signal is estimated as $t_{max} = (t_{32} + t_{23})/2$. (Notice that t_{23} and t_{32} themselves cannot be used as reference points, since their positions depend on the actual signal shape, see Fig. 1.)

For the accurate and fast phase estimation, floating-point arithmetic is only used on the host computer. Therefore the low-end devices preprocess the signal using integer arithmetic, and transmit intermediate results to the host. These intermediate parameters $DSUM$, SUM and NUM are derived as follows:

$NUM = n$ is the number of the reference points in the record. First the device calculates the sum of the distances of the consecutive reference points:

$$DSUM = \sum_{i=2}^n P_{ref}(i) - P_{ref}(i-1) = P_{ref}(n) - P_{ref}(1) \quad (1)$$

where $P_{ref}(i)$ is the distance between the i -th maximum point and the beginning of the measurement. From these intermediate results the wavelength can be calculated on the host, using floating point operations as follows:

$$L = \frac{DSUM}{NUM-1} \quad (2)$$

The RSSI signal's phase is also measured using the reference points. A simple phase can be estimated as follows:

$$\vartheta = 2\pi \frac{P_{ref}(1)}{L} \quad (3)$$

where $P_{ref}(1)$ is the estimate of reference time delay (RTD). To decrease the effect of noise in the estimate of RTD, all of the $P_{ref}(i)$ values are used in the phase estimator, as follows:

$$\vartheta = 2\pi \frac{l}{L} \quad (4)$$

where l is a composite estimate of RTD.

Notice that the following equation holds:

$$P_{ref}(i) = P_{ref}(1) + L(i-1) \quad (5)$$

Let us define parameter SUM as the sum of the distances between the reference points and the beginning of the signal, as follows:

$$SUM = \sum_{i=1}^n P_{ref}(i) \quad (6)$$

Using (5) and (6), the composite estimate of RTD can be calculated as follows:

$$l = \frac{(SUM - \frac{(NUM-1) \cdot NUM \cdot L}{2})}{NUM} \quad (7)$$

which can be simplified to:

$$l = \frac{SUM}{NUM} - \frac{DSUM}{2} \quad (8)$$

The estimated phase value can be calculated from the intermediate parameters at the host computer, using (8) and (4). From two phase values the phase difference can be calculated as follows:

$$\Delta\vartheta = (\vartheta_1 - \vartheta_2) \pmod{2\pi} \quad (9)$$

The operation of the algorithm is illustrated in a real measurement example where radio interferometric signals were generated with two transmitters, and the RSSI signals were measured and processed on two receiver nodes. All the nodes were placed to fixed positions, thus the phase difference should be constant. The distribution of the measurements can be seen in Fig. 3, where result of 1500 distributed phase measurements are shown, using the proposed algorithm.

III. CONCLUSION

A novel algorithm was proposed which is able to measure the phase difference between the RSSI signals of two radio receivers. The proposed method uses preprocessing on low-end devices with integer arithmetic, and post processing on a host computer, using floating-point arithmetic. This solution allows efficient utilization of scarce resources of low-end devices, allowing fast processing required in tracking applications [2]. The proposed algorithm increases accuracy using multiple reference points and averaging.

REFERENCES

- [1] M. Maroti et al, "Radio Interferometric Geolocation," In SENSYS 2005, San Diego, CA, pp. 1-12, Nov 2-4, 2005.
- [2] G. Zachár and G. Simon, "Radio-Interferometric Object Trajectory Estimation," in Proc. Third International Conference on Sensor Networks (SENSORNETS 2014), Lisbon, Portugal, Jan. 7-9, 2014
- [3] B.J.Dil and P.J.M.Havinga: „Stochastic Radio Interferometric Positioning in the 2.4 GHz Range,”. In SENSYS 2011, Seattle, WA, pp. 108-120, Nov 1-4, 2011

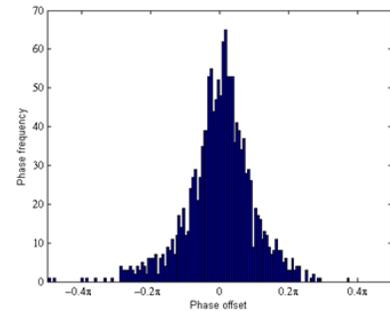


Fig. 3. Distribution of measured phase differences, using the 868 MHz ISM radio-band in disturbance-free environment

Poster abstract: Robust Sensor Fusion for Radio Interferometric Tracking

Gergely Zachár, Gyula Simon
 University of Pannonia, Veszprém, Hungary
 e-mail: {zachar, simon}@dcs.uni-pannon.hu

Abstract—A new sensor fusion is proposed, which provides robust radio interferometry-based object tracking. The proposed solution utilizes redundant measurements and can efficiently handle noisy, erroneous, and missing sensor readings. The efficiency of the proposed fusion is illustrated by simulations and real measurement results.

I. INTRODUCTION

An RF based positioning solution was proposed using radio interferometry [1]. A faster radio interferometric method was recently proposed which can be used for object tracking [2]. The tracking solution uses fixed transmitter and receiver infrastructure nodes to track the movement of a receiver node, when its initial position is known. Now a novel sensor fusion method is introduced, which provides enhanced and more robust position estimation for radio interferometric tracking. The proposed algorithm utilizes redundant measurements to enhance the accuracy of the position estimator and it is also able to self-correct measurement errors.

II. PREVIOUS WORK

Radio Interferometric Positioning [[1]] was proposed to provide low cost localization in wireless sensor networks, utilizing inexpensive off-the-shelf radio components. The method relies on radio interferometry and performs low-frequency signal processing of the interference signals, instead of high frequency signal processing. The method utilizes two transmitter nodes (A and B) to create an interference signal at two receiver nodes (C and D), as shown in Fig. 1. If the frequencies of the two transmitters are approximately the same ($f_A \approx f_B \approx f$) then the interference signal at the receivers will have a low frequency envelope, where the beat frequency is equal to the difference of the two transmit frequencies. This low frequency signal is the actual RSSI, which is available on most RF transceivers.

By measuring the phase difference ϑ between the RSSI signals detected at two receivers (C, D), a metric on the relative positions of the four nodes can be computed as follows [[1]]:

$$\vartheta(f) = 2\pi \frac{d_{AD} - d_{BD} + d_{BC} - d_{AC}}{c/f} \pmod{2\pi} \quad (1)$$

This research was supported by the Hungarian Government and the European Union and co-financed by the European Social Fund under project TÁMOP-4.2.2.C-11/1/KONV-2012-0004. Gergely Zachár was supported by the European Union and co-financed by the European Social Fund in the framework of TÁMOP 4.2.4. A/2-11-1-2012-0001 'National Excellence Program'.

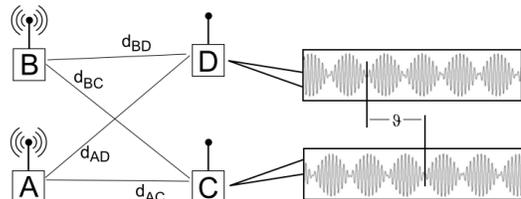


Fig. 1. Radio interferometric phase measurement

where c is the speed of light, and see Fig. 1 for distance notations. Note that this measurement does not provide conventional ranging, because (i) in (1) linear combination of unknown ranges is used, and (ii) only a 2π modulus of the phase is available in (1). The result of (i) is that possible solutions are on a hyperboloid surface, while the effect of phase ambiguity (ii) is that the solution will be a set of hyperboloid surfaces. The phase ambiguity can be resolved either by using multiple carrier frequencies and multiple $\vartheta(f)$ values, resulting rather slow processing [1], or by repeated measurements with phase unwrapping (using only one frequency) [2]. The proposed solution provides a more robust and fault tolerant fusion for the approach of [2].

III. PROPOSED SOLUTION

A. Measurement

Multiple fixed transceiver infrastructure nodes with known positions are used to track a moving transceiver node. In each measurement three infrastructure nodes and the tracked node is used, where two of the infrastructure nodes play the role of transmitters and one infrastructure node and a tracked node play the role of receivers, according to Fig. 1. Thus each measurement provides a phase difference value ϑ in the range of $[0, 2\pi]$. Using repeated measurements with the phase unwrapping technique proposed in [2], the solution of (1) results a single hyperboloid surface for each measurement.

In a *measurement round* multiple (n) measurements are carried out by changing the roles between the infrastructure nodes. Thus in each measurement round a set of independent solutions can be obtained in the form of n hyperboloids. Ideally the intersection of these hyperboloids provides the actual position estimator of the tracked object. However, due to phase measurement errors, the hyperboloids may not cross exactly the ideal position: (i) phase noise (in the range of $[0, \pi]$) causes small and random-like offset, while possible unwrapping errors of [2] (corresponding to phase error of $2k\pi, k = \pm 1, \pm 2, \dots$) may cause large and permanent outliers. The proposed fusion can handle both effects.

B. Basic location estimation

For sake of simplicity the method is discussed in 2D. In each measurement round n hyperbolas are calculated. The following error function is defined for a given point (x, y) :

$$err(x, y) = \sum_{i=1}^n d_i(x, y), \quad (2)$$

where d_i is the distance of point (x, y) from the i^{th} hyperbola. The estimated position (x_e, y_e) is defined as the point where the error function is minimal:

$$(x_e, y_e) = \operatorname{argmin}_{x,y} err(x, y) \quad (3)$$

Although the basic location estimator (3) performs somehow better than the estimator in [0], it still suffers from the effect of large and permanent phase unwrapping errors.

C. Self-correcting position estimation algorithm

Since the calculated hyperbolas, due to phase unwrapping errors, can be far from the true tracked position, the estimator of (3) can have large bias. Note that (2) can be used to detect such situations. Assuming that the number of incorrect phase measurements containing unwrapping errors is much smaller than that of the good ones, the following algorithm can detect and correct these errors.

The idea of the algorithm is the following: in each measurement round an initial position estimator is calculated using (3). Then the possible outlier hyperboloids are tried to be compensated by $\pm 2\pi$, and (3) is reevaluated. If the error (2) is smaller after the trial, the associated phase value is permanently corrected, and the corrected value is used in the next measurement round.

The implemented process uses some restrictions to increase accuracy and to speed up the process. The maximum number of hyperbolas that can be moved is bounded by *check_limit*. The error limit, below what no further enhancement is tried, is defined by *cost_limit*. The pseudo code of the algorithm is shown in Fig. 2. Helper function *calc_error(P)* first calculates (x_e, y_e) from phase measurement set P , using (3), then calculates $d_i(x_e, y_e)$ and $cost = err(x_e, y_e)$, using (2), while *estimate_position* calculates (x_e, y_e) , using (3).

```

P0 = {p1 .. pn} // set of unwrapped phase measurements
Hchecked = {} // set of checked hyperbolas
[cost, d1..dn] = calc_error(P0) // see (2), (3)
WHILE |Hchecked| < check_limit &&
    cost < cost_limit DO
    i = argmax1( d1..dn ), where i ∉ Hchecked
    P+ = (P0 \ pi) ∪ pi + 2π
    P- = (P0 \ pi) ∪ pi - 2π
    CHOSE Pnew ∈ {P0, P+, P-}
    FOR WHICH cost IS MINIMAL, WHERE
    [ncost, nd1..ndn] = calc_error(Pnew)//see (2), (3)
    cost = ncost
    d1 = nd1, d2 = nd2, ... , dn = ndn
    P0 = Pnew
    Hchecked = Hchecked ∪ i
END
[xe, ye] = estimate_position(P0) // see (3)

```

Fig. 2. Pseudo code of the self-correcting position estimation algorithm

IV. EVALUATION

The performance of the self-correcting position estimation algorithm was tested both using simulations and real measurements.

During the simulations a node was moved on a fixed trajectory and the ideal phase measurements were calculated using (1). To the ideal phase values we added normally distributed random noise with variance $\sigma = 0.3\pi$. The noisy phase measurements were processed by the unwrapping algorithm in [2] and then by the proposed self-correcting algorithm. In the test 6 nodes (and correspondingly, 15 hyperbolas) were used for tracking the moving object. The true and estimated trajectories of the object are shown in Fig. 3(a), with blue and green colors, respectively. For the same input the output of method [2] is shown for comparison (in red).

A real measurement result is shown in Fig. 3(b), where a node, carried by a person, was tracked by 4 infrastructure nodes, using 6 phase measurements in each round.

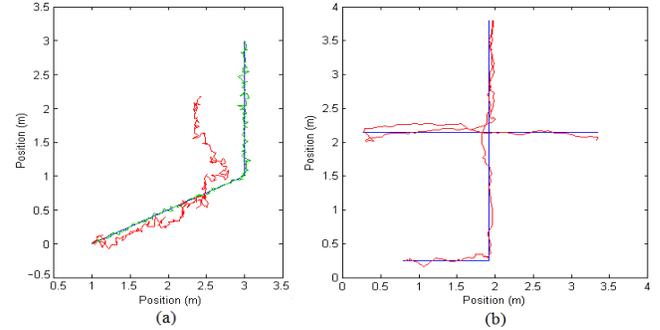


Fig. 3. Simulation and real measurement results. (a) tracking a simulated object with phase noise variance $\sigma=0.3\pi$ and using 6 infrastructure nodes. Blue line: ideal trajectory, green line: proposed method, red line: method of [2]. (b) real time tracking of a moving person. Blue: ideal trajectory, red: proposed method

V. CONCLUSION

A new sensor fusion algorithm was proposed for radio interferometric tracking. The algorithm utilizes redundant measurements to enhance the accuracy of the position estimator and it is also able to correct phase unwrapping errors, often present in real measurements. A novel positioning and self-correcting algorithm was introduced, the performance of which was evaluated using simulations and a real tracking experiment.

According to tests redundant measurements increased robustness and accuracy of the position estimate. Real measurement tests with 4 measurement nodes (and 6 phase measurements in each round) showed that the accuracy of the estimator was around 0.04m in average, with a maximum error of 0.12m.

REFERENCES

- [1] M. Maroti et al, "Radio Interferometric Geolocation", ACM Third International Conference on Embedded Networked Sensor Systems (SenSys 05), San Diego, CA, pp. 1--12, November, 2005.
- G. Zachár and G. Simon, "Radio Interferometric Object Tracking". Technical report. [Online]. Available: <http://www.dcs.uni-pannon.hu/~simon/publications/TechRepRT2013.pdf>

Poster Abstract: Delivering Intelligent Home Energy Management with Autonomous Agents

Conor Muldoon and Michael J. O’Grady and David Lillis and Tadhg O’Sullivan and
Thomas Holz and Gregory M.P. O’Hare

CLARITY: Centre for Sensor Web Technologies, School of Computer Science and Informatics,
University College Dublin

Abstract—This poster discusses the design and implementation of the decision making/reasoning infrastructure of an intelligent home energy management system that was developed as part of the Autonomic Home Area Network Infrastructure (AUTHENTIC) Project. Specifically, the poster focuses on the Agent Factory Micro Edition (AFME) functionality that enables the Home Area Network (HAN) to be managed for two home energy management scenarios representative of this space. The energy management system was tested and deployed in both laboratory and real home settings.

I. INTELLIGENT HOME ENERGY MANAGEMENT

In this poster, the decision making infrastructure of an intelligent home energy management system that was developed using Agent Factory Micro Edition (AFME) [2] is discussed. To drive the development of the home energy management system, two use cases were considered and tested in rooms within the Tyndall National Institute, University College Cork and a real home. The rooms were equipped with temperature, light, contact, and PIR sensors along with pressure mats, which were placed on seats.

In the first use case, when the occupant enters the room, their presence is detected and light levels are adjusted to their preferred settings. If the temperature is below a point specified by the occupant, the air heater is switched on. If the temperature is too high, a message is displayed on the occupant’s tablet device to indicate that the room is being cooled. When the occupant sits down in front of the TV, chair occupancy is detected by a pressure mat. The light level is then adjusted and the TV is switched on through the use of a smart plug. When the occupant leaves the chair, light levels are switched back and a warning message is displayed on the TV for a manual switch off. When the occupant leaves the room, both the air heater and light are switched off. At any stage, if the window is opened, it is detected by a contact sensor and a message is displayed on the occupant’s tablet.

In the second use case, a weather event is simulated through the use of a Waspnote that can detect water. When the event is detected, a message is displayed on the occupant’s tablet and the water heater is turned on, which is simulated by turning the kettle on through, again the use of a smart plug. The idea is that the system should adapt to occupant behaviour – heating water for a shower when the occupant is on their way home from work and it is raining for instance. When the water reaches the occupant’s preferred temperature, availability is shown to

the occupant on the tablet. If the water is not used within a certain time period, an alert message is displayed. If water is repeatedly wasted, the occupant is asked to change their heating schedule. Subsequently, an energy and cost report is displayed on the occupant’s tablet.

The reasoning/decision making module to realise the use cases has been implemented with AFME. AFME is a minimized footprint intelligent agent platform for the rapid development of Multi-Agent Systems. It is based on the Agent Factory development framework [1], but designed for use with the Java Micro Edition (JME) Constrained Limited Device Configuration (CLDC). Although primarily intended for highly constrained devices, applications developed for JME CLDC can also be used on desktop and server machines.

AFME is concerned with the development of computationally reflective agents. Computational reflection is a technique that enables a system to maintain meta-information about itself (an agent’s belief set) and to use this information to determine its behaviour. The behaviour of agents in AFME is represented using declarative antecedent-consequence rules, referred to as commitment rules, that determine the conditions under which commitments are adopted and actions are performed. To facilitate this, the conditions are matched against the agents’ belief sets at periodic intervals using resolution-based reasoning. Resolution-based reasoning is the goal-based querying mechanism that is employed within Prolog interpreters. The reasoning process results in either failure or in a set of bindings being identified that cause commitment rules to be evaluated as true, leading to a number of commitments being adopted and actions being performed.

The AUTHENTIC Reasoning Module comprises a set of agents, a set of software actuators, a set of perceptors, and the AUTHENTIC Service. The AUTHENTIC Service is a class that enables agents to interact with SIXTH, which is a Java-based middleware for the Sensor Web [3] that allows sensor-driven applications to be abstracted from the sensors they depend on. It provides a unified interface that enables a variety of sensor types to be integrated along with a standardised way for interacting with them. SIXTH allows the behaviour of sensors and physical actuators to be altered through the use of a re-tasking service. It is a modular framework and facilitates component updates in a distributed manner and without the need for a restart. For example, if the agent code were updated by the developer and the HAN was in operation in a number

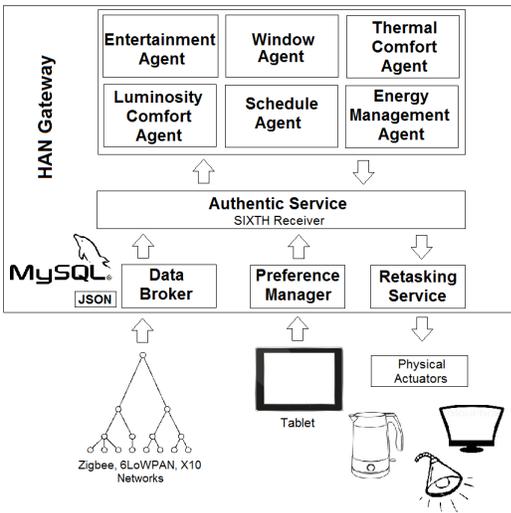


Fig. 1. Data flow within the HAN

of homes, the updated code could be deployed without having to perform on-site installations or restarts.

At present, within the AUTHENTIC Reasoning Module, there are 6 agents (see Figure 1). The Thermal Comfort Agent ensures that temperature levels in the room match the user preferences. If the temperature in the room is at a lower level than the occupant has specified, it sends a command to the HAN to turn the air heater on. Conversely, if the temperature is too high, a command is sent to cool the room. At present, this results in a cooling message being displayed to the occupant. The Luminosity Comfort Agent acts in a similar manner to the Thermal Comfort Agent, but with regard to light – changing the lighting to the appropriate levels using a lamp and a smart plug. The Window Agent informs the occupant when a window is opened by displaying a message on the tablet and the Entertainment Agent controls the application behaviour when the occupant is sitting down and watching television. This results in the light dimming in cases where the light level is high. The Energy Management Agent proactively heats water in anticipation of occupant behaviour. In cases whereby the water has been wasted, the Schedule Agent informs the occupant with an alert message. Subsequently, it requests that heating preferences be changed.

As can be seen from Figure 1, the agent-based reasoning and decision making is performed on the HAN Gateway (base station), which is an Intel Mini ATX. There are two reasons for this: (1) the HAN makes use of a pre-existing communications middleware that has been designed for sensing without regard to device actuation and (2) the types of sensors used are not Java-based and, thus, cannot host AFME agents. The communications middleware delivers messages to SIXTH in the form of JSON messages, which are converted to a standard SIXTH format and delivered to the database along with the AUTHENTIC Service, which also receives information from the Preference Manager. In relation to information coming

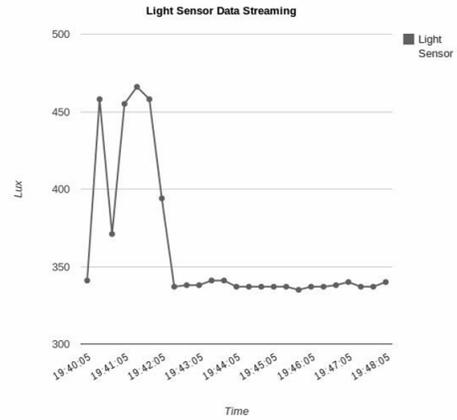


Fig. 2. Graph of streaming data from light sensor

the Data Broker, AUTHENTIC provides Google Analytics graphing functionality (see Figure 2). Perceptors, which are Java classes, are used to enable agents to obtain various type of information received by the AUTHENTIC Service. Agents make decisions with regard to this information and trigger software actuators that cause commands to be sent to physical actuators via the AUTHENTIC and re-tasking services. At present, the user must explicitly specify their preferences in terms of their desired lighting levels, heating levels, etc. Future work will investigate the use of (1) Machine Learning algorithms to enable the system to learn user preferences and needs over time in an autonomous manner and (2) in-network decision making to control devices within the HAN.

II. CONCLUSION

This poster presented the decision making/reasoning functionality of the AUTHENTIC home energy management system. The functionality was delivered through the implementation of a group of AFME agents. The system was tested using two home energy management use cases and deployed both within laboratory and home settings. Future work will investigate the use of Machine Learning algorithms to enable the system to autonomously learn user preferences along with in-network device actuation.

ACKNOWLEDGMENTS

The authors would like to acknowledge the support of the AUTHENTIC Consortium, the International Energy Research Centre under grant TC2012100A, and the Science Foundation Ireland under grant 07/CE/I1147.

REFERENCES

- [1] Rem Collier, Gregory O'Hare, Terry Lowen, and Colm Rooney. Beyond prototyping in the factory of agents. In *Multi-Agent Systems and Applications III*, pages 383–393. Springer, 2003.
- [2] Conor Muldoon, Gregory M. P. O'Hare, Rem W. Collier, and Michael J. O'Grady. Agent factory micro edition: A framework for ambient applications. In *Computational Science–ICCS 2006*. Springer, 2006.
- [3] Gregory M. P. O'Hare, Conor Muldoon, Michael J O'Grady, Rem W Collier, Olga Murdoch, and Dominic Carr. Sensor web interaction. *International Journal on Artificial Intelligence Tools*, 21(02), 2012.

Poster abstract: experimental analysis of environmental perturbations on wireless sensor network operation

Mustafa Mousa, Ahmad Dehwah and Christian Claudel
 Department of Electrical Engineering and Computer Sciences, KAUST, Saudi Arabia
 christian.claudel@kaust.edu.sa

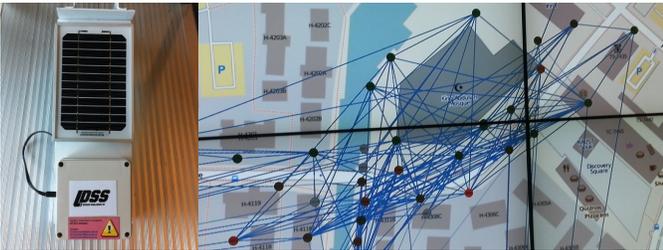


Fig. 1. One unit mote (left) and sensor network visualization (right)

Abstract—Wireless sensor networks are one of the most promising technologies for monitoring large scale environments such as smart cities. However, as all physical systems, their performance is affected by several environmental parameters. In this article, we present the results of a 2-month experiment involving 30+ wireless sensor nodes deployed in an urban area around KAUST campus. Multivariate regression is used to infer the most significant environmental parameters and their effects, and the implications of these results for optimal energy management in sensor networks is discussed.

I. INTRODUCTION

Sensing large scale environments can be achieved through *wireless sensor networks* (WSNs), which consist in a wireless mesh of sensing nodes. A large number of applications of WSNs are in outdoor environments, in which the environmental parameters have considerable variability, unlike indoor environments that are designed for human habitation. Being physical systems, wireless sensor networks are impacted by these environmental parameters, which can negatively affect their performance. Quantifying this performance degradation is very important in practice, as network control schemes and energy management schemes have to take these effects into account.

II. SENSOR NODES

For this study, we used commercial Libelium Waspnotes, which are derived from the Arduino platform series. It is built around a ATmega 1281 microcontroller, with 8 kB of RAM and 128 kB of programmable Flash ROM.

The experiment involved 34 nodes and started by late February 2013, lasting for two months. To study the relation of the

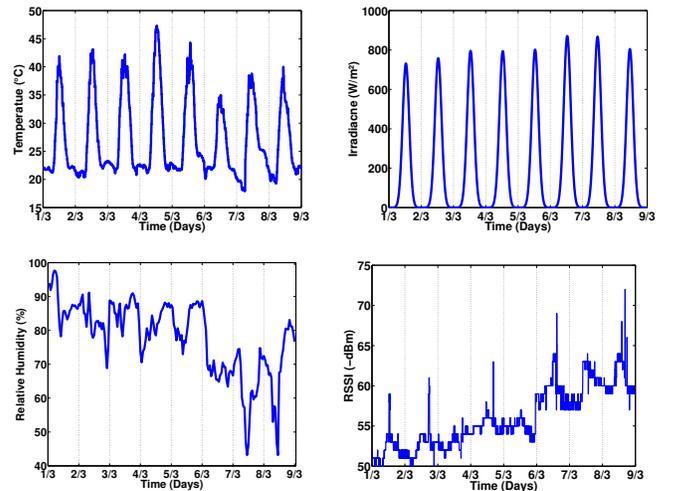


Fig. 2. Temperature (top left), irradiance (top right) and relative humidity (bottom left) measured by the weather station. RSSI average timeseries for a representative link (bottom right)

temperature and irradiance effect on RSSI and PPR we decided to do this study over 8 consecutive days (between two gateway failures) starting March 1st. In this dataset temperature and humidity both have short term periodic (daily) and long-term (weekly) variations. We use as the ground truth data the data generated by the local solar power testing station, which generates solar panel temperature, irradiance and humidity measurements every 900s. The ground truth meteorological data is illustrated in Figure 2.

III. LINK PERFORMANCE ANALYSIS

A. Impact of humidity

There is significant debate in the literature on the actual effects of humidity and temperature on link quality in sensor networks. Thelen et al. [3] conclude that a higher relative humidity improves the received, while Anastasi et al. [1] report that rain and fog cause a decrease in packet reception ratios. Interestingly, these findings contradict the fact that radio signals on frequencies below 11 GHz should be unaffected by rain and fog, as they are transparent to water.

To distinguish the individual effects of the environmental parameters on link quality, we did two multivariate regressions on both link RSSIs and packet reception ratios between nodes, and show the results in Figure 3

B. Impact of temperature

While we can measure the solar panel temperature accurately, the mote temperature can significantly differ from the solar temperature due to air cooling and to internal energy dissipation. In first approximation (neglecting blackbody radiation), the temperature increase of the enclosure of area S , thickness d and thermal conductivity κ subjected to a power forcing P is $\delta T = \frac{P \cdot d}{\kappa \cdot S}$. In the present case, the motes (which dissipate 0.2W) are 0.1C hotter than the panel temperature (which is equal to the air temperature) during the night, and can be up to 7 degrees cooler than the panel temperature with an (irradiance) of $800W/m^2$ (when they are in shadowed areas).

The effect of the temperature on link quality is shown in Figure 3 below.

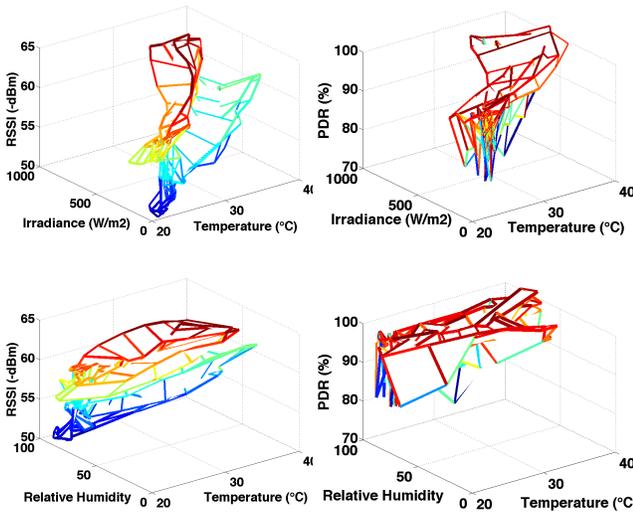


Fig. 3. **Environmental effects on link quality.** Top left: RSSI versus panel temperature and irradiance. Top right: packet delivery ratio versus panel temperature and irradiance. Bottom left: RSSI versus panel temperature and humidity. Bottom right: packet delivery ratio versus panel temperature and humidity.

As can be seen from the above Figure, temperature has the largest impact on link quality. In fact, a multivariate regression on temperature, irradiance and humidity gives the following results:

$RSSI = \alpha \cdot T + \beta \cdot I + \gamma \cdot H$, where $\alpha = 95$, $\beta = -35$ and $\gamma = 11$. The uncertainties (from the covariance matrix) are $\Delta\alpha = 9$, $\Delta\beta = 4$ and $\Delta\gamma = 5$. Thus the RSSI is very strongly correlated with panel temperature, negatively correlated with irradiance and statistically uncorrelated with humidity. The negative correlation with irradiance is expected as air temperature is positively correlated with panel temperature and negatively correlated with irradiance.

Temperature affects link quality since it increases the level of thermal noise (which in turn reduces the probability of successful decoding).

IV. POWER ANALYSIS

A. Power consumption

Since our motes are not currently equipped with battery charging current sensors, we have to estimate their consumption by analyzing the battery voltage drop over time during the night (when no solar power forcing is present). Since the voltage-charge relationship in a battery is dependent on the battery condition and discharge rate, we estimate the battery condition using the procedure outlined in [2]. Results are shown in Figure 4 below.

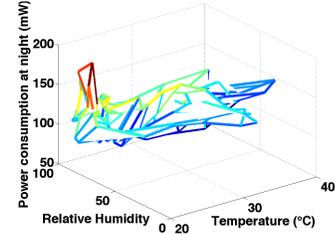


Fig. 4. **Impact of environmental effects on estimated night power consumption.** Power drawn from the battery versus air temperature and humidity.

A similar multivariate regression gives the following results: $Power = \alpha \cdot T + \beta \cdot H$, where $\alpha = 0.5$ and $\beta = 3$. This shows the limits of multivariate linear regression, as our data is far from being normally distributed in the present case. Thus the dependency with the humidity is mainly driven by a few outliers for extreme humidity values. Power dissipation is approximately constant over the 40 – 90% humidity range.

Power dissipation increases with temperature since the resistivity of semiconductors decrease with temperature, causing higher current drain (at the same voltage). The outliers appearing in extreme humidity conditions could be caused by condensation of water on the electronics in some motes due to enclosure damage.

V. FUTURE WORK

Our objective is to investigate a new mote designed by our group This mote can sense more environmental parameters, including its own temperature, its battery temperature, the battery charging or discharging current and the solar power input. This will allow us to identify more accurately the effects of environmental parameters on wireless sensor network operation.

REFERENCES

- [1] G. Anastasi, M. Conti, M. Di Francesco, and V. Neri. Reliability and energy efficiency in multi-hop ieee 802.15.4/zigbee wireless sensor networks. In *Computers and Communications (ISCC), 2010 IEEE Symposium on*, pages 336–341, 2010.
- [2] Mustafa Mousa and Christian Claudel. Energy parameter estimation in solar powered wireless sensor networks. In *Proceedings of the 2013 REALWSN conference*.
- [3] John Thelen, Dann Goense, and Koen Langendoen. Radio wave propagation in potato fields. In *1st Workshop on Wireless Network Measurements*, volume 2, pages 331–338. Citeseer, 2005.

Poster abstract: Intelligent Building Management Systems Based on Cognitive M2M Communications

Manolis Surligas*, Stefanos Papadakis*, Vangelis Angelakis†, Sofoklis Kyriazakos‡, Elias Tragos* and Di Yuan†

*Institute of Computer Science, Foundation for Research and Technology-Hellas, GR-71110 Heraklion, Crete, Greece.

‡Converge ICT Solutions & Services, 74 Panormou Str, GR-11523 Athens, Greece.

†Department of Science and Technology, Linköping University, SE-601 74 Norrköping, Sweden.

Corresponding author: surligas@ics.forth.gr

Abstract—Intelligent “green buildings rely on automated energy management systems to reduce energy consumption and operational costs. This is typically accomplished by a large number of building-wide deployed smart objects, sensing the environment and acting accordingly. Existing systems are mainly centralized, introducing availability, scalability and robustness issues, having a single point of failure and relying on human interaction for their actions. In this work we propose a highly distributed, self-organizing, self-managing, smart objects system for building management, based on wireless communication, that reaps benefits from its cognitive, distributed and adaptive design.

I. INTRODUCTION

Buildings hold a large share in the worldwide energy consumption. Commercial and residential buildings were responsible for 41% of the total energy consumption in 2010 in the US [1] and for 40% in 2009 in the EU [2]. Aiming to reduce carbon emissions, the European Commission has published in June 2011 a directive for reducing energy consumption of 20% by the year 2020 [3]. One step towards achieving this ambitious goal is to equip buildings with intelligence, enabling them to operate autonomously and efficiently with minimum human interaction, take decisions and act without affecting the habits and the well-being of the tenants.

The focus of the intelligent buildings is on providing automated functionalities i.e. for energy management, heating, ventilation, air-conditioning (HVAC) and access control. To achieve these, a large number of indoor and outdoor smart objects (sensors and actuators) should be deployed ubiquitously in the building. These smart objects are capable of sensing a variety of environmental indicators, i.e. temperature, luminance, wind speed, humidity, CO₂ levels, motion and human presence and act according to user preferences. The “brain” of an intelligent building is the Building Management System (BMS), which is responsible for the robust, reliable and efficient coordination of the deployed smart objects, as well as of their decisions and actions.

Existing BMS solutions utilize predefined rules for coordinating the sensing and actuating devices, functioning typically in a centralized way (see Figure 1(a)) with a central unit receiving input from all devices and taking the final decisions for i.e. when to open the air conditioning. Furthermore, human intervention is required for adding new devices or re-configuring the already existing ones, introducing the possibility of human error. These systems impose severe limitations not only by the

single point of failure, but also by the small scaling capability imposed by the human factor [4], [5].

For reducing the installation cost of a BMS wireless technology is utilized for connecting the smart objects. Wireless Radio Frequency (RF) signals can easily propagate through walls, so there is no need for a special cabling infrastructure, making these systems also applicable to older buildings. Furthermore, RF communication is less vulnerable in emergency and extreme conditions, i.e. a fire. For coordination and communication of this kind of networks, several standards and network architectures have been developed, i.e. IEEE 802.15.4 [6], Zigbee [7], 6LoWPAN [8] and WirelessHART [9].

Various standards and communication protocols are used among vendors, which restricts the interoperability of the devices and requires the installation of excess devices, wasting resources, while increasing complexity and cost. In case of a device failure, the system may become unstable, take wrong decisions or become unable to communicate with some of the devices, due to the incompatibility of devices of different vendors. Commonly, wireless smart objects utilize the unlicensed ISM bands to transmit information introducing many challenges [10] as these bands are excessively crowded by high data rate transmissions. On the other hand, smart objects may affect in turn the WLAN in terms of backoff time [10]. Therefore, the increasing number of smart objects may lead to performance degradation of the “primary” spectrum users.

II. SORBET APPROACH

Aiming to address the drawbacks of existing BMS solutions, SORBeT (*Smart Objects for Intelligent Building Management*) will develop an architectural framework enabling the highly distributed, self-organizing, self-managing, fault-tolerant interconnectivity of smart-objects for the robust management of intelligent “green” buildings. SORBeT is a new Marie Curie Industry-Academia Partnerships and Pathways (MC-IAPP) FP7 project which focuses on enhancing system reliability, ensuring continuous and self-healing connectivity utilizing different communication technologies, while keeping performance and energy costs at competitive levels. The methodology adopted by SORBeT to achieve these goals is split into four steps: (i) analysis of the requirements of intelligent buildings in terms of the input smart objects need to provide and the applications enabled within them, (ii) development of a framework to interconnect large numbers of heterogeneous smart objects using virtualization and Cognitive Radio (CR)

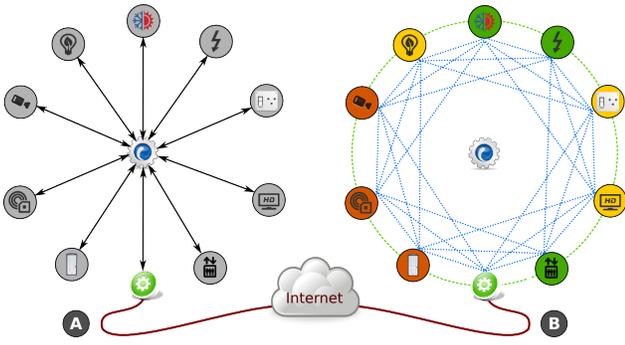


Fig. 1: SOrBeT approach

technology, (iii) modeling of the trustworthiness of the smart objects to enable reliable bootstrapping and self-monitoring of the network to detect and mitigate failures, and (iv) implementation of a modular open software platform to demonstrate the performance of the developed schemes.

SOrBet tries to minimize not only the human interaction but also the dependency on central coordinating nodes by promoting the M2M (machine-to-machine) communication of the smart objects (see Figure 1(b)), giving the possibility to organise the objects into clusters (depicted with the different colors in Figure 1(b)). That way, smart objects will perform a series of tasks, take decisions and act locally without the need of a central decision unit. This approach improves the reliability of the system because any local incidents/failures will not affect the overall system and contributes to the scalability of the system due to the fact that all the decisions are taken locally by a limited group of objects.

Towards avoiding wireless interference and improving the reliability and availability of smart objects, SOrBeT aims to utilize the advantages of CR technology. Several solutions have been proposed (i.e. [11], [12]) to reduce the effects of interference in the overcrowded ISM bands, but they are not always efficient, especially when the smart objects have to transmit lots of data (i.e. a camera transmitting video for remote surveillance). Based on the fact that many licensed bands are underutilized, they could be used for the sensors' wireless communication, without affecting primary transmissions, i.e. the sub-900 MHz band that allows better electromagnetic wave propagation even through thick walls in a building. Studies have shown that taking advantage of CR technology, the sensors' performance and energy consumption can be improved drastically [13]. Thus, combining an M2M approach with the cognitive capabilities of the smart objects, the potential of SOrBeT is exponentially widened.

Another strong feature of our approach, is the way decisions are taken. Information regarding "trust" both from the user and the smart object perspective, is taken into consideration. This can be achieved by constructing a reputation mechanism along with a user behavior tracking system that identifies, flags and avoids users or smart objects behaviors that could have negative effect to the system performance. This feature shields our system from potential jamming attacks or malicious users [14]. Along with the reputation mechanism a strong security design will be applied vertically in the system in order to provide further immunity from vulnerabilities al-

lowing data integrity, confidentiality and source identification.

III. CONCLUSION AND EXPECTED RESULTS

In current BMS systems there are emerging issues of robustness, reliability, availability and scalability, mainly due to their centralized nature and the single point of failure. SOrBet expects to address these issues, providing a distributed, self-organized and adaptive BMS platform, using state of the art technologies and techniques. Smart objects cooperate and act without the need of central coordinators in an M2M style of communication. More precisely, our approach provides reliable, self-healing, plug and play connectivity of a large number of smart objects. Due to its cognitive capabilities, our approach will set the basis for a new era in wireless sensor networks, providing a road-map towards CR enabled smart buildings. In general, SOrBeT will develop a holistic, advanced, reliable and trustworthy platform for building automation, with a specific focus on monitoring and minimizing the energy consumption of buildings. From a small house to a large building, SOrBet achieves to satisfy the need for an advanced, secure and efficient system, enabling services like remote management, access control, physical intrusion or fire detection, all fully optimized in terms of performance and reliability.

ACKNOWLEDGMENT

This work has been supported in part by the EC Marie Curie Actions project SOrBeT (FP7-PEOPLE-2013-IAPP: 612361)

REFERENCES

- [1] U.S Department of Energy. Buildings energy data book.
- [2] M. Economidou *et al.*, "Europe's buildings under the microscope," *A Country-by-country review of the energy performance of buildings*, 2011.
- [3] "Com 2011 370 directive of the european parliament and of the council on energy efficiency and repealing directives 2004/8/ec and 2006/32/ec."
- [4] A. Mitseva *et al.*, "Intelligent system for independent living and self-care of seniors with mild cognitive impairment or mild dementia," *Journal on Information Technology in Healthcare*, vol. 7, no. 6, 2009.
- [5] O. Evangelatos, K. Samarasinghe, and J. Rolim, "Evaluating design approaches for smart building systems," in *IEEE MASS Workshop on Internet of Things Technology and Architectures*, 2012.
- [6] "IEEE Standard for Local and metropolitan area networks-Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)," *IEEE Std. 802.15.4-2011*, 2011.
- [7] ZigBee Alliance, "Specification document 053474r17," vol. 1, 2008.
- [8] "Transmission of IPv6 Packets over IEEE 802.15.4 Networks," Internet Requests for Comments, RFC 6282, September 2007.
- [9] Hart communication foundation. [Online]. Available: <http://www.hartcomm.org/index.html>
- [10] C.-J. M. Liang *et al.*, "Surviving wi-fi interference in low power zigbee networks," in *Proceedings of the 8th ACM SenSys*. ACM, 2010.
- [11] W. Xu *et al.*, "Channel surfing: defending wireless sensor networks from interference," in *Proceedings of the 6th IPSN*. ACM, 2007.
- [12] E. Tragos *et al.*, "Spectrum assignment in cognitive radio networks: A comprehensive survey," *IEEE Commun. Surveys Tuts.*, vol. 15, 2013.
- [13] E. Tragos and V. Angelakis, "Cognitive Radio Inspired M2M Communications (Invited Paper)," in *Proc. IEEE Global Wireless Summit 2013*.
- [14] W. Xu, K. Ma, W. Trappe, and Y. Zhang, "Jamming sensor networks: attack and defense strategies," *Network, IEEE*, vol. 20, no. 3, 2006.

Poster Abstract: Secure Metering Communication in the Energy Smart Home Lab

Anton Hergenröder
Institute of Telematics
Karlsruhe Institute of Technology
Email: hergenroeder@kit.edu

Kaibin Bao
Institute of Applied Informatics
and Formal Description Methods
Karlsruhe Institute of Technology
Email: bao@kit.edu

Thomas Bräuchle
Center for Applied Legal Studies
Karlsruhe Institute of Technology
Email: thomas.braeuchle@kit.edu

Abstract—*Smart Metering* is an important step in the evolution of home automation. With comprehensive data gathering of smart metering systems, there is also demand for security. Furthermore, smart homes participate in the *internet of everything*, our highly interconnected world, which raises even more security challenges. In this work we present our work-in-progress on investigating security aspects of smart metering systems in smart homes.

I. INTRODUCTION

Responsible usage of electrical energy is becoming increasingly important as Germany's government set an objective to rely on renewable energy sources by at least 80 percent in 2050 [2, § 1]. Also European and G8 leaders announced in 2008 the goal to reduce greenhouse gas emissions by at least 80 percent below emission levels of 1990 [5]. These goals require the electrical energy supply and utilization to be optimized for efficiency and comfort. Modernization of the public power supply comes into existence as so called *Smart Grid*. Controllability of energy consumption is the second most important requirement for a Smart Grid. We believe the climate goals can only be reached in the future, if domestic buildings are equipped with an energy management infrastructure, which monitors the energy consumption of household appliances and controls their operation, driven by signals from the Smart Grid. Doing so, it is possible to optimize the energy consumption and reduce the overall costs, by running appliances during times of high power availability (i.e. during high solar radiation or spanking breezes). This is convenient for appliances like a dish washer, a washing machine or an electric car. The integration of households with domestic intelligent energy management infrastructures (*Smart Homes*) into the Smart Grid begins with *Smart Metering* systems, which monitor the usage of electrical energy, heating and water supply and are capable of communicate the collected data to the commodity providers and/or metering service providers. In Figure 1 the described Smart Home scenario with it's components is illustrated.

II. CHALLENGES AND REQUIREMENTS FOR SECURITY

Significant challenges emerge during the development of energy management systems for smart homes. At first, many sensors are needed to monitor the appliances and all relevant environmental information. Temperature, humidity, lighting

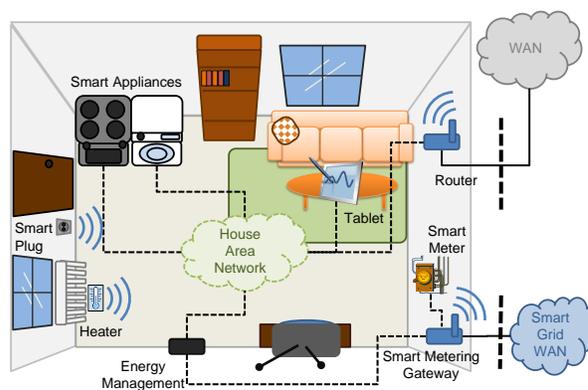


Figure 1. Smart Metering in Smart Home Scenario

conditions of each room in a household must be collected as well as energy consumption and status information of all relevant appliances. Since it is very costly and complicated to connect all the sensors by wire, it is convenient to use wireless sensor networks (WSNs). Second, there is demand to information security. Collecting all the sensor data in a smart home, an attacker could learn personal information of the inhabitant. Manipulation of the sensor data or control messages can also lead to safety risks, e.g. when driving a heater too hot or remotely control the stove. Therefore, the Federal Office of Information Security (BSI) in Germany has published a protection profile and technical guidelines for the communication gateway of a smart metering system [6]. Consequently, the security requirements defined in this document should also be met by the whole smart home.

In the scope of the project *KASTEL* [1], we investigate security aspects of future smart home solutions in an interdisciplinary researcher team. One of our goals is to extend an existing living lab architecture with a secure smart metering solution based on wireless sensor networks. The existing living lab called Energy Smart Home Lab [3][4] is a real standalone smart home building on the campus of Karlsruhe Institute of Technology (KIT). In our work, we address the challenges

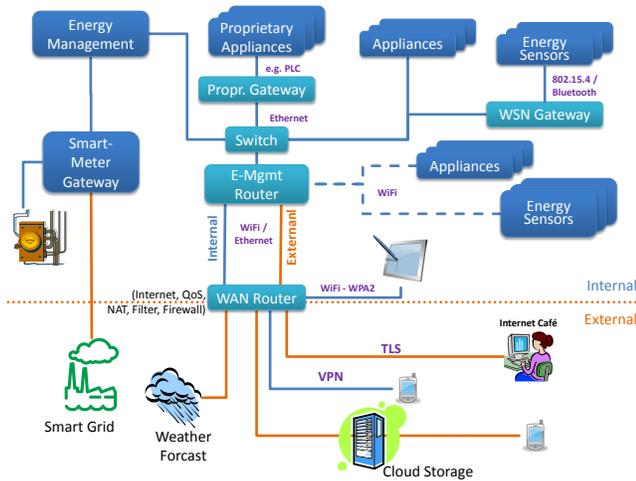


Figure 2. Integration of Smart Metering Network into a Smart Home

of secure communication in smart homes, including WSNs. Sensor nodes have to be small and cheap, which results in resource restrictions typical for WSNs, namely little energy and low computational power. Therefore communication and security mechanisms have to be energy-efficient. For energy-efficient security we use hardware-based security modules attached to standard sensor nodes to compute cryptographic functions efficiently and fast [7].

Figure 2 shows the network architecture of the secure smart home. There is the smart meter and its gateway (as described in [6]), which is connected to a service provider outside the household. The gateway also delivers metering and tariff information to the household energy management system. The energy management system itself is connected to all internal household appliances and sensors. An inhabitant of the household has the ability to access all the collected data of the energy management by various means and paths. To support different access profiles which depends on the access path, a supporting e-management-router has the ability to determine wherever data is accessed from inside or outside of the household.

III. LEGAL ASPECTS OF SMART HOME SECURITY

In order to integrate lawful and intelligent Smart Metering Systems in the energy sector, binding regulations for data protection and IT-Security have to be taken into account. As the collection, provision and storage of meter data from one or more Smart Meters of one or multiple commodities entails the risk of determining a highly resolved profile of energy consumption, it is mandatory to implement data protection and IT-Security regulations [8]. Therefore the Federal Office of Information Security (BSI) has published a Protection Profile [6] and a Technical Guideline for the Gateway of a Smart Metering System to ensure a level of protection according to the state of the art.

Those documents define the IT-security objectives and requirements for Smart Metering Systems with a Smart Meter

Gateway as the central communication component. Furthermore the Protection Profile and the Technical Guideline is aimed at developers of Smart Meter Gateways and contains the binding requirements that have to be implemented. The security functionality of the Smart Meter Gateway contains provisions dealing with the protection of confidentiality, authenticity, integrity of data and information flow control mainly in the purpose of protecting the privacy of consumers, to ensure a reliable billing process and to protect the Smart Metering System [6, p. 7].

In the Smart Home scenario, a Smart Meter Gateway is a part of an intelligent energy management system inside the Home Area Network (HAN). It could be used e.g. to monitor the energy consumption of household appliances. In addition the energy management is able to collect, process and store more data than just the records from the Smart Meters, e.g. temperature and humidity data as well as the energy consumption. For this reason it is necessary to analyze all information flows inside the energy management system of a Smart Home consisting of the Smart Meter Gateway and a multitude of sensors sending relevant environmental information about the consumer.

IV. ACKNOWLEDGMENTS

This work is partially funded by the German Federal Ministry of Education and Research.

REFERENCES

- [1] Center of Excellence for Applied Security Technology (KASTEL), German: Kompetenzzentrum für angewandte Sicherheitstechnologie. <https://www.kastel.kit.edu/>.
- [2] *German Renewable Energy Law (EEG)*. 2009.
- [3] F. Allerding and H. Schmeck. Organic Smart Home: architecture for energy management in intelligent buildings. In *Proceedings of the 2011 workshop on Organic computing*, pages 67–76, 2011.
- [4] B. Becker, F. Allerding, U. Reiner, M. Kahl, U. Richter, D. Pathmaperuma, H. Schmeck, and T. Leibfried. Decentralized energy-management to control smart-home architectures. In *Architecture of Computing Systems-ARCS 2010*, pages 150–161. 2010.
- [5] EUROPEAN COMMISSION. Energy roadmap 2050, December 2011.
- [6] German Federal Office for Information Security. Profile for the Gateway of a Smart Metering System (SMGW-PP), March 2013.
- [7] A. Hergenröder and C. Haas. Energy-efficient security in smart metering scenarios. Demo on the 38th IEEE Conference on Local Computer Networks (LCN), October 2013. Sydney, Australia.
- [8] O. Raabe, M. Lorenz, F. Pallas, and E. Weis. Harmonisierung konträrer Kommunikationsmodelle im Datenschutzkonzept des EnWG. In *Computer & Recht*, volume 27, pages 831–840, 2011.

Demo Abstract: Activity Rate Adaptation in WSN

Victor Cionca, Wensi Wang and Brendan O’Flynn
 Wireless Sensing Group, Tyndall National Institute, Cork, Ireland
 victor.cionca@tyndall.ie

Abstract—Wireless Sensor Network (WSN) nodes relying on energy harvesting must maintain the system energy supply between a low and a high threshold, to prevent energy exhaustion and waste. Node activity is mostly periodical, with short bursts of power consumption. By controlling the activity period, the energy consumption can be increased, or reduced. This work presents a scheduling algorithm for solar-powered WSN. The algorithm relies on a prediction of expected harvested energy for the next day to decide the optimal activity rate in the current time slot. The algorithm uses a combination of long and short term planning, as well as proportional control for dynamic rate adaptation.

I. INTRODUCTION

As energy harvesting technology is maturing, it is becoming an attractive solution for powering WSN. Energy harvesting WSN don’t use the generated energy directly but store it into a buffer, such as super capacitors or thin-film batteries. Solar power is the most common ambient energy source, partly due to its regular patterns that are relatively easy to predict. Most WSN nodes support a low-power sleep mode, and most of the functionality of nodes is periodic and duty-cycled: the node wakes up from low-power mode, performs the activity at a power level P_a , then goes back to sleep. Energy harvesting creates the need for power management, and activity rate adaptation: as the harvested energy varies, a constant activity rate would either lead to energy waste or energy overspend and exhaustion, both undesirable. Current solutions see this as an optimization problem, where the object is to maximise duty cycle [1] throughout time slots, while maintaining the energy supply constant; other approaches use control theory [2] or time-constrained scheduling [3].

We present a simple activity rate management algorithm tailored for solar-powered systems. The algorithm takes advantage of the increase-decrease pattern observed in harvested solar power, to make sure that the system energy supply lies between pre-defined minimum and maximum thresholds. The algorithm builds upon simple proportional control with a short-term and long-term planning scheme that relies on prediction of expected harvested energy.

II. RATE ADAPTATION ALGORITHM

In our investigation of activity rate adaptation, we started with a proportional control algorithm, which attempts to maintain the system’s energy level constant by varying the active time. The error between the current system energy and its set point is proportional to the amount of time the system spends in active mode. The linear relation $err = \alpha \cdot T_a + \beta$

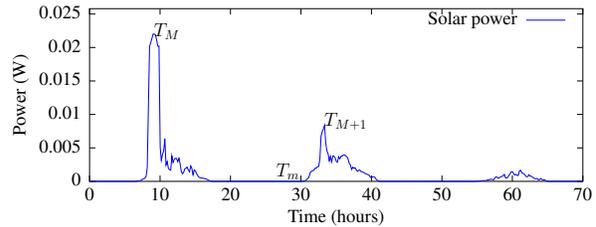


Fig. 1. Solar power trace over three days. Highlighted on the graph are the extreme points, that mark the beginning of the increase or decrease slopes.

can be fitted via linear regression, and at runtime the system active time is computed as $T_a = (err - \beta)/\alpha$.

As shown in figure 1, solar power follows a day-night cycle, with a maximum point towards the middle of the day and a low point at night time, as well as increasing and decreasing slopes. Proportional control computes the optimal activity rate to maintain the system’s energy constant, when solar power increases. However, the algorithm does not plan ahead: when solar power decreases, it can use too much energy, preventing it from reaching the set point on the next increase slope, if the energy harvested then is too low.

To improve the algorithm, a prediction of the expected harvested energy is required. Considering that the solar power curve has the characteristics shown in figure 1 (a minimum and a maximum points, increase and decrease slopes), the algorithm’s goal is to make sure that:

- at the minimum point of the solar power curve, the system energy is greater than the minimum threshold
- at the maximum point, the system energy is equal to the maximum threshold.

To achieve this, the algorithm employs *short-term* and *long-term* planning. In the *short-term* planning, the algorithm runs the proportional control algorithm at the start of each slot; however, it calculates the active time such that it will have saved a certain amount of energy, Δ_s , at the end of the slot. The *long-term* planning determines the value of Δ_s to meet the goal specified above, as follows:

$$\Delta_s = \frac{E_{target} - E_{crt}}{t_{target} - t_{crt}}, \quad (1)$$

where t_{target} is the next time slot when the harvested energy will reach an extreme (minimum or maximum), and E_{target} is the corresponding limitation for system energy. The value of Δ_s needs to be computed in every slot, but determining t_{target} , which is a search operation, must only be done when the system has reached an extreme.

The worst case scenario is that when the solar power increases, the algorithm must be able to reach E_{max} with zero activity, by staying in sleep mode for the whole duration. If the energy harvested during the increase slope, E_{gen} , is too low, the pre-defined threshold, E_{min}^{pre} might be too low, and E_{min} must be re-calculated:

$$E_{min} = \max(E_{min}^{pre}, E_{min}^{real}), \quad (2)$$

$$E_{min}^{real} = E_{max} - E_{gen} + E_{sleep}. \quad (3)$$

The calculation must be done each time the system reaches the high extreme.

To summarise, the algorithm works as follows:

- 1) Determine the current slope of the solar power.
- 2) Determine the next extreme, t_{target} and E_{target} ; if the solar power starts decreasing, update E_{min} as per eq. 2.
- 3) At the beginning of each slot:
 - calculate Δ_s using eq. 1
 - determine active time, $T_a = (\Delta_s - \alpha)/\beta$
 - distribute T_a to all system activities
- 4) Upon reaching the extreme, go to step 2.

III. ANALYSIS

The algorithm has been tested in a simulated environment, with simulated solar power as well as with real data acquired from light intensity traces and extrapolated into harvested power. To predict the expected solar power an EWMA (Exponentially Weighted Moving Average) was used. The simulation models a real system developed in-house, with a maximum power consumption of $0.1W$, a sleep power of $20\mu W$ and an application that consumes $0.074W$ for $20ms$ each time it is active. Figure 2 shows the results of a simulated run, using real values for harvested solar energy. Our system uses a $2.5F$, $5V$, super-capacitor, therefore high energy threshold is set at $30J$ and the low threshold at $20J$. The traces of harvested solar power span ten days, and the simulated system is in active mode 1.2% of the time, just over 3 hours. The average overspent energy per slot is $0.097J$. The wasted energy is much larger, because the generated energy much surpasses the consumed energy.

The algorithm's goal is for the system to start every decrease-increase cycle at the high energy threshold, which would ensure perpetual functionality. With perfect prediction of expected harvested energy, the algorithm will adapt the activity rate on the decreasing slope, to save the necessary amount of energy that will allow it to reach the high threshold again at the end of the increasing slope. Of course, this requires a perfect predictor, which does not exist. The EWMA predictor used in this work is not good at predicting sudden changes, and behaves *conservatively* upon under-estimation (predicts less energy, therefore system saves more than necessary), and *optimistically* upon over-estimation (predicts more energy, therefore system consumes more than necessary). We tried to enhance the accuracy with online weather forecast, as proposed in [4] but found the online weather data too coarse

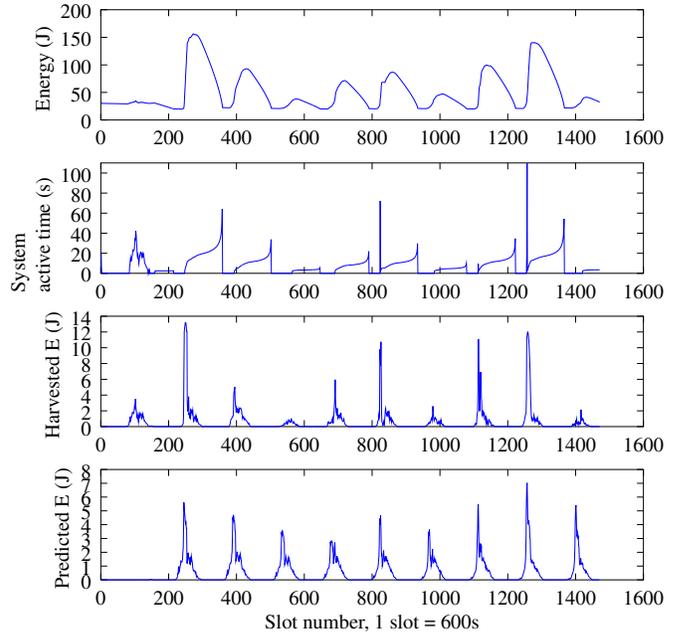


Fig. 2. System energy level and active time determined by the algorithm with real harvested energy data

to be useful. More complex algorithms are available, but their requirements are prohibitive for the WSN environment.

The algorithm has been implemented in TinyOS and this demonstration will show, through an additional graphical interface, how the activity rate is varied to maintain the energy constant, as a lamp simulates a day-night cycle.

IV. CONCLUSIONS

This paper has presented an algorithm to dynamically adapt a system's activity rate in order to maintain the system's energy between a minimum and maximum thresholds. The algorithm relies on the periodical nature of solar power, with increasing and decreasing slopes, as well as being able to predict the amount of energy generated in a certain interval.

ACKNOWLEDGEMENTS

This work was funded by Enterprise Ireland and the International Energy Research Center (IERC) as part of the AUTHENTIC project.

REFERENCES

- [1] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava, "Power management in energy harvesting sensor networks," *ACM Trans. Embed. Comput. Syst.*, vol. 6, no. 4, sep 2007.
- [2] C. M. Vigorito, D. Ganesan, and A. G. Barto, "Adaptive control of duty cycling in energy-harvesting wireless sensor networks," in *Sensor, Mesh and Ad Hoc Communications and Networks, 2007. SECON'07, 2007*, pp. 21–30.
- [3] C. Moser, D. Brunelli, L. Thiele, and L. Benini, "Real-time scheduling for energy harvesting sensor nodes," *Real-Time Systems*, vol. 37, no. 3, pp. 233–260, 2007.
- [4] N. Sharma, J. Gummeson, D. Irwin, and P. Shenoy, "Cloudy computing: Leveraging weather forecasts in energy harvesting sensor systems," in *Sensor Mesh and Ad Hoc Communications and Networks, 2010. SECON'10, 2010*, pp. 1–9.

Poster Abstract: Sustainable Backpressure Data Collection in Solar Powered Wireless Sensor Networks

Shusen Yang and Julie A. McCann

Department of Computing, Imperial College London, UK
Intel Collaborative Research Institute (ICRI) for Sustainable Connected Cities

{s.yang09, j.mccann}@imperial.ac.uk

Abstract—Solar Powered WSNs (SP-WSNs) have been attracting a growing interest in various research fields including hardware system design, power management schemes, and network-wide algorithms. Renewable solar energy provides an opportunity to achieve Energy Neutral Operation (ENO) (nodes always have energy), while posing new challenges for designing energy-adaptive network algorithm. In this poster, we propose a energy-aware routing protocol SP-BCP, which can achieve sustainable data collection for environmental monitoring applications in SP-WSNs. We evaluate the performance of SP-BCP through simulations and experiments on a real SP-WSN platform.

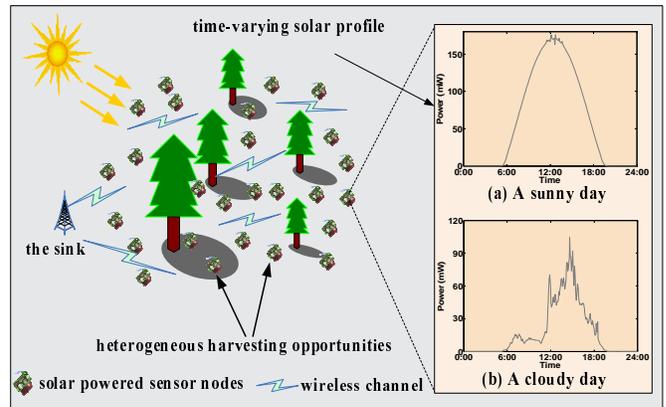


Fig. 1. Illustration of a typical multi-hop SP-WSN.

I. INTRODUCTION

Harvesting energy (e.g. solar, thermal, wind and vibrational) from the environment brings a step change that ensures the viability of Wireless Sensor Networks (WSNs) for real world deployments. Compared with other environmental energy sources, solar energy has been more widely considered and Solar Powered WSNs (SP-WSNs) have been attracting a growing interest in various research fields. However, due to the limited sizes of micro solar panels, harvested solar energy remains scarce [1], [2]. To make best use of this resource one needs to understanding the dynamics of energy generation. Further complications come about because of the heterogeneous spatial harvesting capabilities across different nodes in a sensing space due to shading or cloud coverage as shown in Figure 1.

These have significant impacts on the network algorithm design in SP-WSNs. For instance, a routing protocol should adaptively select and dynamically adjust end-to-end paths to avoid the time-varying routing hot-spots (i.e. nodes with low harvesting power). In this poster, we develop an energy-adaptive routing protocol, SP-BCP, by applying a throughput-optimal backpressure routing to SP-WSNs. We show that SP-BCP can achieve energy neutral operations (ENO) [3], through testbed experiments and simulations.

II. SOLAR POWERED BACKPRESSURE COLLECTION PROTOCOL

Routing is one of the most important functionality in multi-hop WSNs at the network layer. We consider the backpressure routing that is well-known for dynamically routing data packets over multi-hop networks by using nodes' local information (e.g. queue backlog and link quality), based on the Lyapunov drift minimization theory. Recently, a practical routing protocol, BCP [4], has been developed for data collection in WSNs. BCP is adaptive to sudden link fluctuations, queue hot-spots, and topology changes. However, without energy awareness, BCP can not provide any guarantee for efficient solar energy usage or long-term ENO in SP-WSNs. Therefore, we develop SP-BCP, an energy aware backpressure routing protocol for sustainable data collections in SP-WSNs.

To ensure ENO, each node should not consume more energy than the energy budget provided by power management component (e.g. [1], [3], [6]). Let the remaining energy of node x be E_x . Let N_x be the set of x ' all one-hop neighbors. For a node x , let the energy costs by transmitting and receiving be ET_x and ER_x respectively. Let Q_x be the queue backlog of a node x and $c_{x,y}$ be the capacity of wireless link (x, y) .

The operations of SP-BCP for packet transmission are described as follows:

- 1) **Energy aware weight calculation:** Each sensor node x computes the weights $w_{x,y}$ for all its neighbors $y \in N_x$. If remaining energy budgets of x and its neighbor y are larger than receiving and transmitting cost respectively (i.e. $E_x \geq ET_x$ or $E_y \geq ER_y$), it sets $w_{x,y} = (Q_x - Q_y)c_{x,y}$ otherwise $w_{x,y}$ is set as zero.
- 2) **Routing:** Each node x selects link (x, y^*) with the maximum weight for optimal potential receiver $y^* \in N_x$.
- 3) **Forwarding:** If $w_{x,y^*} > 0$, x forwards the packet to y^* .
- 4) **Remaining activity duration update:** When a packet is transmitted, both the transmitter x and the receiver y^* update their remaining energy budget as $E_x = E_x - ET_x$ and $E_{y^*} = E_{y^*} - ER_{y^*}$.

The communication overhead of SP-BCP is quite light: every node x can dynamically obtain Q_y and DR_y by periodically broadcasting one-hop beacons or using overhearing. Evaluation results presented in Section III demonstrate that SP-BCP achieves hard ENO guarantee, and therefore prevents any node running out of energy. Beside proposed SP-BCP, we can also use $1/\min(E_x, E_y)$ as the penalty over a link (x, y) , since $1/\min(E_x, E_y) \rightarrow \infty$ as $\min(E_x, E_y) \rightarrow 0$; or stabilize virtual queues to ensure ENO.

III. EVALUATION

We implemented SP-BCP in TinyOS 2.1 and evaluated them through our aforementioned SP-WSN platform and the Tossim simulator. All real-world experiments used the on-line measured solar power, while all simulations were based on the public solar database [5].

To evaluate the performance of SP-BCP, we first compared the real-world performance of BCP and SP-BCP in our 16-node SP-WSN for three days. We implemented power management algorithm proposed in [6] to compute the energy budgets for SP-BCP. We set $D_{\min}=10\%$, the initial battery level as $1.2 \text{ KJ} \pm 30\%$ for different nodes, and sensing (packet generation) rate as one packet per two seconds. As shown in Figure 2, for SP-BCP, no nodes run out of energy during the three days. Therefore, ENO were achieved. The sink continuously received data in every slot and achieved relatively high packet delivery ratio (about 76.8–98.9%) during the daytime. The main reasons of packet loss are limited data buffer size and inelastic sensing rates (a flow controller could significantly reduce the packet loss). For BCP, however, 10 nodes died during the second day, leading to network disconnection and significant degradation of network goodput.

Figure 3 shows the Tossim-based simulation results for a random deployed 50-node SP-WSNs for 12 days. Sensing rate of every node was set as one packet per seven seconds, and initial battery level were set as $200 \text{ J} \pm 30\%$ for different nodes randomly. The simulation shows similar results to the testbed experiment, BCP failed quickly in the second day, but SP-BCP achieved sustainable data collection for the whole 12 days. In summary, both experiment and simulation results

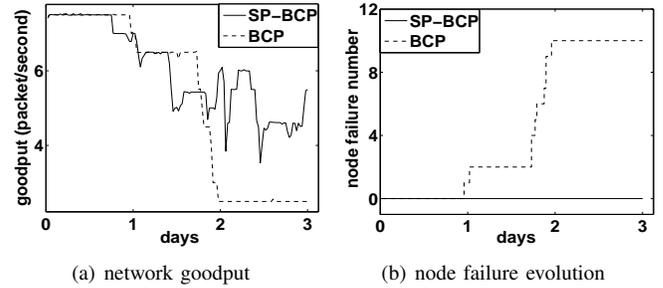


Fig. 2. Real-world experiment results of BCP and SP-BCP.

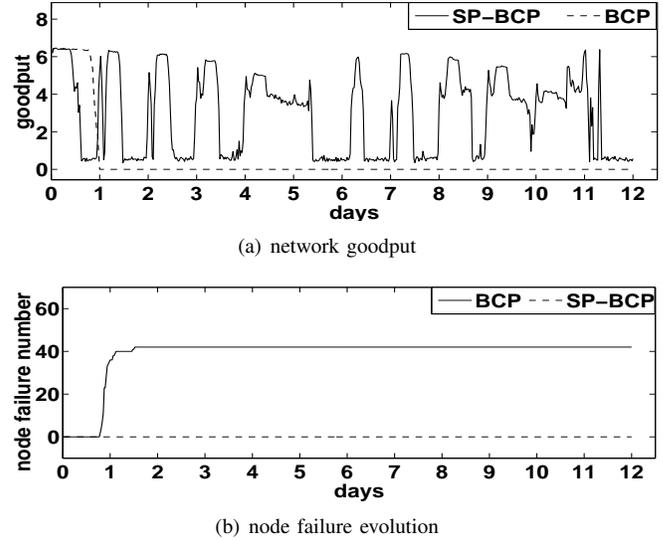


Fig. 3. Simulation results of BCP and SP-BCP.

show that SP-BCP can improve the end-to-end performance of backpressure routing protocols in SP-WSNs.

IV. CONCLUSION AND FUTURE WORK

This poster presents a simple energy aware routing protocol, SP-BCP, for sustainable data collection in solar powered wireless sensor networks. Through real-world experiments and simulations in Tossim, we show that SP-BCP can adaptive to energy dynamics and management to achieve energy neutral operations. Our future work will focus on joint networking and data fusion in SP-WSNs.

REFERENCES

- [1] C. Moser, L. Thiele, D. Brunelli, and L. Benini, "Adaptive power management for environmentally powered systems," *IEEE Trans. Computers*, vol. 59, no. 4, pp. 478–491, 2010.
- [2] S. Sudevalayam and P. Kulkarni, "Energy harvesting sensor nodes: Survey and implications," *IEEE Commun. Surveys Tuts.*, vol. 13, no. 3, pp. 443–461, 2011.
- [3] A. Kansal, J. Hsu, S. Zahedi, and M. Srivastava, "Power management in energy harvesting sensor networks," *ACM Trans. Embedded Comput. Sys.*, vol. 6, no. 4, 2007.
- [4] S. Moeller, A. Sridharan, B. Krishnamachari, and O. Gnawali, "Routing without routes: the backpressure collection protocol," in *Proc. ACM/IEEE IPSN*, 2010.
- [5] <http://solaratd.uoregon.edu/SelectArchival.html>.
- [6] S. Yang, X. Yang, J. McCann, T. Zhang, G. Liu, and Z. Liu, "Autonomic Solar Powered WSN for Network-wide Protocols," *IEEE Journal on selected areas in communications*, in press.

Demo Abstract:

SicsthSense - Dispersing the Cloud

Liam McNamara, Beshr Al Nahas, Simon Duquennoy, Joakim Eriksson and Thiemo Voigt

SICS Swedish ICT AB, Sweden

Email: ljm@sics.se, beshr@sics.se, simondqu@sics.se, joakime@sics.se and thiemo@sics.se

Abstract—This demo presents SicsthSense, our open cloud platform for the Internet of Things. SicsthSense enables low power devices such as sensor nodes and smartphones to easily store their generated data streams in the cloud. This allows the data streams, and their history, to be made permanently available to users for visualisation, processing and sharing. Moving sensor data computation and monitoring into the cloud is a promising avenue to enable centralisation of control and redistribution of collected data.

We showcase SicsthSense running with real sensor nodes collecting environmental data and posting it to our datastore. This live data is then visualised and made available for sharing between users of the platform. Our Android App will also be distributed to enable participants to stream their phone sensors into the system, demonstrating how simple it can be to start machine-to-machine interactions with SicsthSense.

I. MOTIVATION

The “Internet of Things” (IoT) is a unifying term for sensor networks, pervasive computing/connectivity and cloud computing. It represents a large set of modern computing challenges, stemming from the huge scale of the vision. Incorporating billions of devices with a huge variety of heterogeneity in the hardware and applications being run. This idea will involve needs that we cannot even foresee at the moment. Therefore developing solutions with the maximum standards interoperability and flexible extensibility is crucial.

SicsthSense is an open cloud platform for the Internet of Things’ data, designed to ease the interconnection of the billions of sensors and actuators that we expect in the near future. With SicsthSense, devices are registered, discovered, configured, and programmed from the cloud, either through a Web interface or an open machine-to-machine API. This makes it possible to connect heterogeneous devices ranging from smartphones to low-cost embedded sensor and actuators.

We believe the following features are important for promoting easy development of applications in IoT, allowing developers to concentrate on their domain specific problems rather than the “nuts and bolts” of making use of sensed data.

- Easily connect devices to the cloud;
- Store/retrieve sensor data in the cloud;
- Make decisions in the cloud;
- Visualise data;
- Actuate from the cloud.

Solving these problems led to the creation of web-based “resource” repository, where users can have arbitrary de-

vices/websites/URLs periodically polled and parsed to be turned into data streams that are then easy to use in further processing. Alternatively, data can be pushed into the system using our API. All data streams are stored indefinitely, giving users complete historical access. There are also social aspects, whereby users can share reading/writing access to their streams, and subsequently revoke such access.

We aim to facilitate research on the Internet of Things, but also on Big Data analysis, distributed data storage, and interaction design etc. Furthermore, we expect researchers from other fields to benefit from SicsthSense, as users of the platform. SicsthSense supports the HTTP and CoAP protocols, through poll and push interaction schemes, as the main application layer protocols. The system is structured as a true RESTful interface, simplifying the design of client applications. Particular consideration is given for low-power IP devices, such as Contiki nodes using 6LoWPAN and RPL. Figure 1 presents the architecture of SicsthSense.

Long running interactions are enabled by Websockets, so that devices with very high update frequencies can be efficiently supported. Such persistent connections are also used to enable publish/subscribe interactions between devices, so data points can be immediately sent to interested parties, avoiding the need for wasteful polling. Beyond simply storing and retrieving data, SicsthSense also provides features for basic actuation, allowing the system to affect the real world when certain conditions are met by the incoming data streams. For instance, when the average temperature of a user’s house goes below 15C, SicsthSense can signal the heater to switch on.

There are similar commercial offerings in this field, offering a combined data store, visualisation and limited processing engine for IoT applications. The main offering is Xively (formerly COSM), with Open.sen.se and ThingSpeak.com acting as competition. All these commercial products either charge for data access, or require other forms of revenue. SicsthSense not only provides data access free of charge, but allows users to fork the code and run their own instance, giving them complete control of their data and its usage. Other researchers have realised the value of an open academic solution to the problems we have identified [3].

II. ARCHITECTURE

The main SicsthSense platform is a Java-based server that was explicitly designed to be able to run on a wide

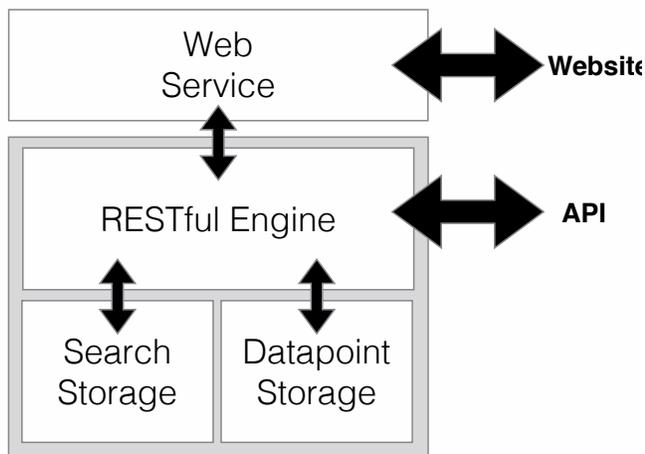


Fig. 1. System Architecture. Showing the main components in the grey box, which provides the machine-to-machine interface, and the optional web interface for a nicer user experience through a web browser.

range of devices, from powerful cloud-based servers to small home gateway machines. The server is distributed as a single .jar file, to enable easy packaging and operation. Storage is provided by two subsystems, first a normal RDBMS that is specifically designed to enable easy searching for relevant feeds. Secondly, a highly scalable asynchronous masterless replication database system is used for storing data points. All the logic, processing and authentication is implemented in the engine sub-system which provides a RESTful HTTP interface for machine-to-machine interactions with SicsthSense. An optional web server is also provided which will provide a user friendly interface to the system. This richer interface is particularly useful non-technical users and allows immediate visualisation of collected data in the SicsthSense system.

Communication with the API is performed using JSON to represent entities. So when users need to request attributes of some part of our system (e.g. user definition or a data point), they simply perform an HTTP GET on a well defined URL and receive a JSON representation of that entity. HTTP POSTs to the same URL can then reconfigure the system.

III. DEMO

The demo will consist of a live demonstration of datastreams being pulled and pushed into the SicsthSense platform. This data will then be visualised, processed and distributed out to endpoints. We will operate a set of sensor nodes running our SicsthSense image of Contiki, each having data posted and/or polled into the SicsthSense system. We will show how the system can be distributed into many federated silos, allowing a local copy of SicsthSense to be run in one's own network, providing low latency and high bandwidth connections. This local silo will then communicate with the wider SicsthSense server, whenever connectivity allows. Our demo will have some sensor nodes communicating with a Raspberry Pi acting as this gateway on a user's own network, before synchronising with the cloud server. The demo setup is illustrated in Figure 2.

Local processing and actuation can also be performed by these local silos, simplifying connectivity and security issues that present themselves when attempting to control devices

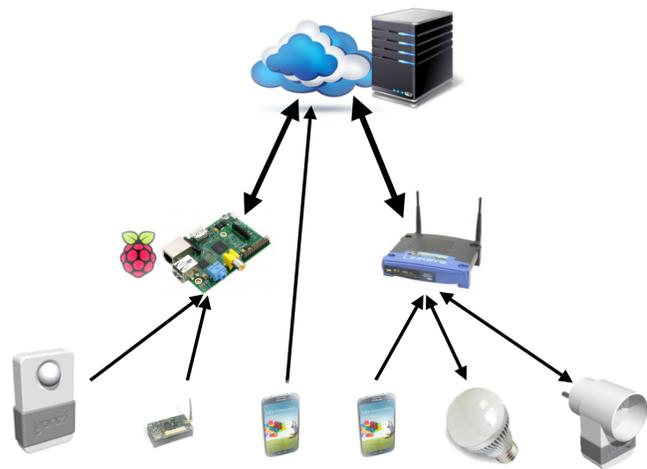


Fig. 2. Demo Setup. With many sensing devices (sensor nodes, IoT devices and phones) all storing their data in the cloud server. Some devices are directly connected, while others connect through a gateway. Some gateway connected devices have their updates aggregated before being transmitted to the cloud.

directly from the cloud. This will be shown through the Raspberry Pi controlling LED lamps that will be present at the demo dependant on the incoming data streams. All the data collected throughout the demonstration will be publically available via the cloud service.

Attendees will be invited to install a SicsthSense App for their Android smartphone to post data themselves. This App will send the phone's current battery life, connectivity strength and other features, allowing their data to be visualised and shared with other attendees at the conference. The simplicity of interaction with our platform will be demonstrated through these interactions, showing how little code and complexity is required to achieve the desired outcomes.

IV. CODE RELEASE

We provide a live running instance of SicsthSense for free (<http://sense.sics.se>) [2] so users can interact with the cloud data platform without having to run their own server. We also release its source code under a BSD open source license [1]. This is of particular use where, for example, a company needs to deploy its own, private, instance of SicsthSense due to sensitivity of their data, or lack of connectivity. Meaning that most importantly the code is open source and free. This repository also contain code for the Contiki image and the android client. It is freely forkable and has already been used as the starting point for other projects. Documentation for creating clients and consumers of the data is also provided on the project Wiki page.

ACKNOWLEDGMENTS

This project was supported by VINNOVA, Energimyn-digheten (the Swedish Energy Agency), and EIT ICT Labs.

REFERENCES

- [1] Available on github here <https://github.com/sics-iot/sicsthSense>.
- [2] Live instance of the platform running at <http://sense.sics.se> and an early feature testing site at <http://presense.sics.se>. 2013.
- [3] Michael Blackstock and Rodger Lea. Iot mashups with the wotkit. In *Internet of Things (IOT), 2012 3rd International Conference on the*, pages 159–166. IEEE, 2012.

Demo Abstract: How Temperature Affects IoT Communication

James Brown and Utz Roedig
School of Computing and Communications
Lancaster University, United Kingdom
{j.brown, u.roedig}@lancaster.ac.uk

Carlo Alberto Boano and Kay Römer
Institute for Technical Informatics
TU Graz, Austria
{cboano, roemer}@tugraz.at

Nicolas Tsiftes
SICS Swedish ICT
Kista, Sweden
nvt@sics.se

Abstract—In the future we will rely on applications built on top of the Internet of Things (IoT). Example applications are smart cities, smart grids and smart healthcare. These IoT applications require a reliable service with predictable quality, and that sensor data and actuator commands are delivered reliably and timely. Unfortunately, IoT performance is highly affected by environmental conditions, especially by ambient temperature. It is therefore necessary to configure an IoT system such that sufficient application performance is provided under all environmental conditions that may be encountered.

In this demonstration, we show how ambient temperature affects the performance of an IoT application. Specifically, we connect remotely to a large-scale temperature-controlled testbed and show how temperature affects the operation of a state-of-the-art routing protocol. Using a local setup, we further demonstrate how the impact of temperature on communication links can be easily captured and modelled in order to inform the design of communication protocols robust to high temperature fluctuations.

I. INTRODUCTION AND MOTIVATION

The Internet of Things (IoT) will enable applications of utmost societal value including smart cities, smart grids and smart healthcare. For the majority of such applications, strict dependability requirements are placed on IoT performance, and sensor data as well as actuator commands must be delivered reliably and timely. Whilst existing IoT solutions are designed to provide such dependable performance, many typically fail, as embedded wireless systems are significantly affected by their often hostile environment. Radio interference from other wireless equipment and electrical appliances impairs communication, whilst temperature and humidity variations affect battery capacity and electronics.

Recent research has revealed that operations of wireless sensor systems are largely affected by their on-board temperature. Temperature variations can significantly affect the quality of wireless links [1], battery capacity and discharge [2], as well as clock drift [3]. Deployed wireless sensor systems can undergo substantial variation in temperature depending on the enclosure and on the deployed location. Systems exposed to sunshine can easily experience high temperatures up to 70 degrees Celsius – especially if the packaging absorbs infra-red (IR) radiation [4]. High temperature may also be accompanied by high variation as seen in [5], which showed that the on-board temperature of a system can vary in an outdoor deployment by as much as 35°C in one hour and 56°C over the course of a day. In previous work, we showed that this extreme variance can reduce the received signal strength (RSS) between communicating nodes by more than 6 dB [5], which

is enough to change the packet reception rate (PRR) of what was a good link from 100% to 0%.

Whilst RSS fluctuations on a link caused by temperature changes have been well studied [1], [5], the effect of these changes on networking protocols have received little attention. A deep analysis of how such effects on individual links affect the operation of networking protocols is necessary to inform the design of more dependable systems. In this demonstration we present an initial investigation into how variations in temperature affect the received signal strength and the behaviour of state-of-the-art IoT routing protocols. We also demonstrate how temperature effects on RSS can be captured and modelled in order to inform the design of robust networking protocols able to mitigate the effects of temperature variations.

II. DEMO DESCRIPTION

Our demo is made up of two parts. In the first part, we will connect to TempLab [6], a temperature-controlled experimental infrastructure with the ability of precisely varying the on-board temperature of wireless sensor nodes, to demonstrate how temperature affects IoT communications. We put specific emphasis on the operation of a state-of-the-art routing protocol. In the second part, we will demonstrate how such temperature effects on communications can be quantified by measuring the signal degradation using a local demonstration made of four Maxfor MTM-CM5000MSP wireless sensor nodes of which two will be heated using infra-red heating lamps.

Remote testbed at TU Graz. For the first part of the demo, we will connect to a live experiment running on our experimental infrastructure at TU Graz. The testbed will be configured to replay pre-recorded individual outdoor temperature traces on 17 Maxfor MTM-CM5000MSP wireless sensor nodes. The nodes will be running Contiki’s IPv6 Routing Protocol for Low-Power and Lossy Networks (ContikiRPL) [7] with each node periodically sending packets to the sink. We will illustrate the impact that temperature variations have on key network metrics such as throughput, delay and lifetime, showing that it is fundamental that temperature effects are taken into account when designing robust networking protocols. Along with the sensor nodes and controllable infra-red lamps, the testbed facility in Graz offers the possibility to observe an experiment using a steerable Web-cam (Figure 1). This facility will be used during the demo to monitor the live experiment.

Measuring temperature effects on communication. Along-side the observation of the remote demo, we will use a small-scale version of our testbed infrastructure at the EWSN

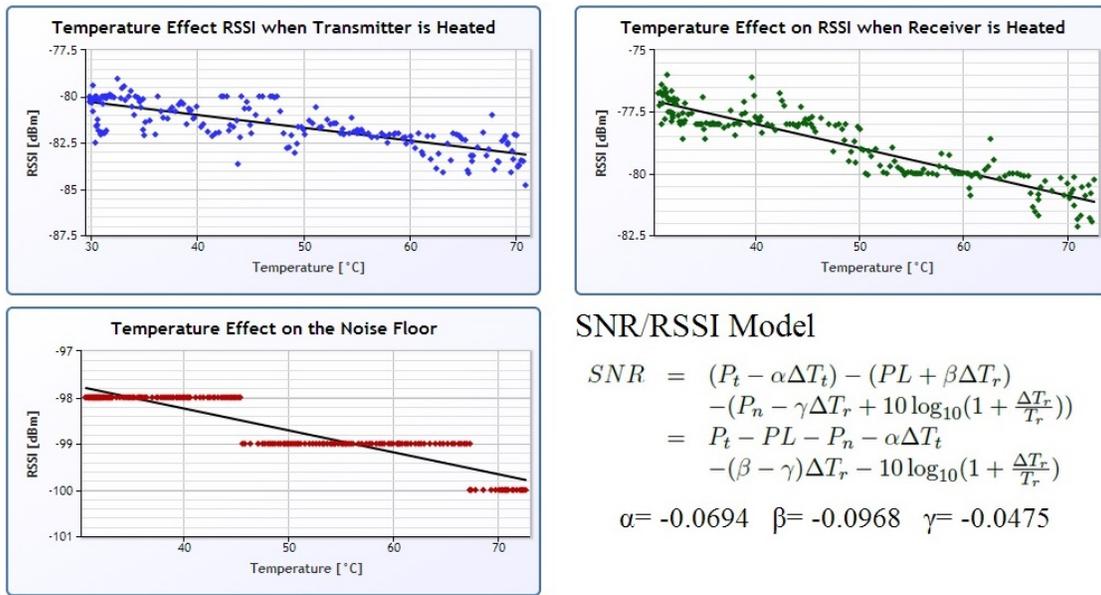


Fig. 2. Real-time capture of the degradation of the wireless link quality as a function of temperature, and parametrization of a first-order model.



Fig. 1. Snapshot of the testbed infrastructure available at TU Graz taken remotely using the steerable Web-cam.

demo session to capture how link quality degrades at high temperatures. Specifically, we will use four Maxfor MTM-CM5000MSP wireless sensor nodes of which two will be heated using infra-red heating lamps similar to the ones shown in Figure 1. The demo will show how the dependency between temperature and link quality can be determined in real-time by measuring temperature and signal strength information.

In particular, we form two independent links by dividing the sensor nodes into pairs. Each link operates on individual channels to avoid effects caused by interference. Sensor nodes will run the same Contiki software periodically sending packets using a predefined transmission power to their intended receiver with temperature information measured using the on-board SHT11 digital sensor. On the receiver, statistics about the received packets¹ are logged using the serial port and collected by an application that computes the decrease in RSS as a function of temperature, as we proposed in [5]. An intuitive web-application will visualise this data as shown in Figure 2.

¹We collect hardware-based link quality metrics in IEEE 802.15.4 compliant radio transceivers, namely the received signal strength indicator upon packet reception (RSSI) and in absence of packet transmissions (noise floor).

III. CONCLUSIONS

Variations in environmental properties can significantly affect the operations of wireless embedded systems, which can break IoT applications. In this demonstration, we have shown that temperature variations can drastically affect the performance of communication protocols, with particular focus on the state-of-the-art ContikiRPL. We have further demonstrated how the temperature effects on received signal strength can be captured to parametrize a first-order model that can be used in future protocols to predict changes in link quality and minimize the impact of temperature fluctuations.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Union 7th Framework Programme (FP7-ICT-2011-8) under grant agreement n° 317826 (RELYonIT).

REFERENCES

- [1] K. Bannister, G. Giorgetti, and S.K.S. Gupta. Wireless sensor networking for hot applications: Effects of temperature on signal strength, data collection and localization. In *Proc. of the 5th HotEmNets*, 2008.
- [2] C. Park, K. Lahiri, and A. Raghunathan. Battery discharge characteristics of wireless sensor nodes: An experimental analysis. In *Proc. of the 2nd SECON Conf.*, 2005.
- [3] T. Schmid. *Time in Wireless Embedded Systems*. PhD thesis, University of California, 2009.
- [4] C.A. Boano, J. Brown, N. Tsiftes, U. Roedig, and T. Voigt. The impact of temperature on outdoor industrial sensor network applications. *IEEE Trans. Ind. Informatics*, 6(3), 2010.
- [5] C.A. Boano, H. Wennerström, M.A. Zúñiga, J. Brown, C. Keppitiyagama, F.J. Oppermann, U. Roedig, L.-Å. Nordén, T. Voigt, and K. Römer. Hot Packets: A systematic evaluation of the effect of temperature on low power wireless transceivers. In *Proc. of the 5th ExtremeCom*, 2013.
- [6] C.A. Boano, M.A. Zúñiga, J. Brown, U. Roedig, C. Keppitiyagama, and K. Römer. TempLab: A testbed infrastructure to study the impact of temperature on wireless sensor networks. In *Proc. of the 13th IPSN Conf.*, 2014.
- [7] N. Tsiftes, J. Eriksson, and A. Dunkels. Low-power wireless IPv6 routing with ContikiRPL. In *Proc. of the 9th IPSN Conf.*, 2010.

Demo Abstract: Testbed Infrastructure for Benchmarking RF-based Indoor Localization Solutions under Controlled Interference

Filip Lemic, Jasper Büsch, Mikolaj Chwalisz, Vlado Handziski, Adam Wolisz
 Telecommunication Networks Group (TKN), Technische Universität Berlin (TUB)
 {lemic, buesch, chwalisz, handziski, wolisz}@tkn.tu-berlin.de

Abstract—The proliferation of RF-based indoor localization solutions raises the need for testing systems that enable objective evaluation of their functional and non functional properties. We introduce a testbed infrastructure for automatized benchmarking of RF-based indoor localization solutions. The infrastructure leverages a robotic mobility platform which serves as a reference localization system and can transport the localized device in an autonomous and repeatable manner. Using a well defined interface, the infrastructure obtains location estimates from the System Under Test which are subsequently processed in a dedicated metrics computation engine. In this demonstration, we present the capabilities of the testbed infrastructure on the use-case of benchmarking a WiFi fingerprinting indoor localization system under the influence of controlled interference.

I. INTRODUCTION

Localization in urban environments and buildings, where people spend most of their time, draws lots of research attention nowadays. Due to the poor indoor performance of the Global Positioning System (GPS) service, a number of alternative localization approaches have been recently proposed. The performance of these solutions is usually evaluated using local ad-hoc experimental setups, mostly by the development teams themselves, using manual and laborious procedures. The results are typically captured following different evaluation scenarios and metrics, with different number and position of evaluation points, and in different environments. Due to these factors, the majority of existing evaluation approaches do not provide sufficient basis for predicting the performance of the systems in other environments and do not enable objective comparison of multiple solutions.

In this work we present a testbed infrastructure for automatized benchmarking of Radio Frequency (RF) based indoor localization solutions that aims at addressing some of these identified shortcomings. It leverages a robotic mobility platform which enables accurate and repeatable positioning of localization device. Furthermore, it integrates devices for generating controlled RF interference, which can be used to evaluate the influence of RF interference on the performance of the localization System Under Test (SUT). For validation of the resulting RF context, the infrastructure features devices that monitor the RF spectrum at different measurement points in order to guarantee equal conditions for all SUTs.

We apply the benchmarking methodology developed within the EVARILOS Project and described in [1]. It follows the

principles of black-box testing and is not privy to the inner operation of the SUTs. The only direct interaction with the SUT, performed over a well defined API, is used to obtain location estimates. Combined with reference location data from the robotic mobility platform, the obtained location estimates are subsequently processed by a dedicated engine that calculates relevant evaluation metrics like geometrical and room level accuracy, latency and power consumption.

In this demonstration we showcase the possibilities of the presented testbed infrastructure on the example of benchmarking a WiFi fingerprinting localization solution with and without controlled interference.

II. TESTBED INFRASTRUCTURE

This section presents the components of our testbed infrastructure for indoor localization benchmarking. The overview of the system is given in Figure 1.

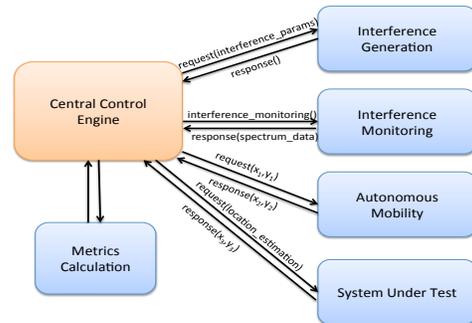


Fig. 1. Overview of the benchmarking system

a) Mobility Support: For automatic transportation of the localized device of the SUT to different evaluation points, without the presence of a human test-person and in a repeatable way, we use the Turtlebot II robotic platform [2]. It comprises of a mobile base called Kobuki, a laptop, a router and a Microsoft Kinect 3D camera sensor. On the software side we are using Robot Operating System (ROS) [3], an open source approach for robots. The Turtlebot can autonomously position itself, together with SUT, at different locations. ROS navigation stack uses the measured depth information from the Kinect sensor and an a-priori given floor plan to localize itself.

Furthermore, it provides an interface to request the robot to drive autonomously to a given coordinate. For an experiment we define a set of measurement points and the robot iterates over each of them. When a measurement point is reached, our central control engine sends a request for location estimation to the SUT. Once the SUT reports the localization results, the robot moves to the next measurement point.

b) Interference Generation: During an evaluation of the proposed localization solution the impact of external interference is mostly not considered. However, it can have an influence on the performance of the SUT. For this reason we have developed means to generate various types of interference scenarios. The most common type of wireless activity in the 2.4 GHz is the WiFi traffic. We have adapted the interference scenarios from [4], [5] and using the cControl and Management Framework (OMF) [6] we are able to create in our testbed the interference context of typical home or office environments.

c) Interference Monitoring: In the previous section we have described that we can generate different interference scenarios based on the needs of an experiment. Still, the spectrum in the ISM 2.4 GHz band is free for use and we do not have full control over all devices operating in those frequencies. We have disabled the university infrastructure WiFi network in the 2.4 GHz band in the building but the signal from the surroundings can still be received. That is why it is necessary to monitor the spectral environment to tell if it is looking as expected. We can use OMF to orchestrate Wi-Spy [7] devices to perform spectrum sensing. We are using one of them connected to the robot to make sure that the measured interference is not exceeding planned one.

d) Location estimation requests and metrics calculation: One part of our testbed infrastructure is an engine for requesting a location estimation from the SUT. Based on the accurate location obtained from the robotic platform and the location estimation obtained from the SUT, the engine calculates the performance metrics. The set of benchmarking metrics that can be calculated or estimated consists of geometrical and room level localization accuracy, latency of location estimation and the power consumption of the SUT. Geometrical accuracy is the distance between the accurate location reported by the robotic platform and the one reported by the SUT. The room level accuracy is a binary metrics containing the correctness of the estimated room. Furthermore, the latency of location estimation is the time that SUT needs in order to report a location estimation. Finally, the power consumption is the power used by the SUT for location estimation, assuming that SUT exposes interfaces to access this information.

III. DEMO DESCRIPTION

In the demo we use our testbed infrastructure to show two experiments remotely over the Internet. The hardware components used for this demo are given in Figure 2. Namely, we show the usage of our indoor localization benchmarking infrastructure with a WiFi fingerprinting localization system as the example of SUT. The control of the experiment is done in a centralized way, using the central control engine. First, we

demonstrate the usage of our infrastructure for benchmarking in the environment without controlled interference. The coordinates of the evaluation points are sequentially given to the robotic platform. When the point is reached, the robot reports its current location to the central engine. Central engine then requests the location estimation from the SUT. When the estimated location is reported from the SUT, the central engine calculates the evaluation metrics. The whole experiment is monitored for a RF interference. Secondly, we repeat the whole experiment with the presence of an artificially created WiFi interference.

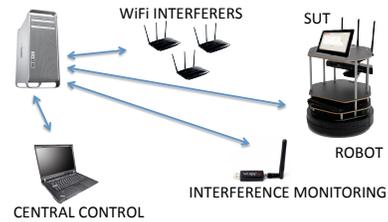


Fig. 2. Hardware components of the infrastructure

IV. CONCLUSION

In this demonstration we introduce a testbed infrastructure for objective benchmarking of RF-based indoor localization solutions. We showcase how four important evaluation metrics can be captured in entirely automated and repeatable way, under two evaluation scenarios, one without and one with exposure to controlled WiFi interference.

V. ACKNOWLEDGEMENTS

This work has been partially funded by the European Commission (FP7-ICT-FIRE) within the projects EVARILOS (grant No. 317989) and CREW (grant No. 258301) and EIT ICT Labs, as part of the activity 12149, “From WSN Testbeds to CPS Testbeds”.

REFERENCES

- [1] V. Haute *et al.*, “D2.1 Initial Version of the EVARILOS Benchmarking Handbook,” 2013.
- [2] *Turtlebot II Documentation*, 2013. [Online]. Available: <http://wiki.ros.org/Robots/TurtleBot>.
- [3] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, “ROS: an open-source Robot Operating System,” in *ICRA workshop*, 2009.
- [4] W. Liu, M. Mehari, S. Bouckaert, L. Tytgat, I. Moerman, and P. Demeester, “Demo Abstract: A Proof of Concept Implementation for Cognitive Wireless Sensor Network on a Large-scale Wireless Testbed,” in *EWSN’13*, 2013.
- [5] *CREW repository*, 2013. [Online]. Available: <http://www.crew-project.eu/repository/>.
- [6] *OMF 6 Documentation*, 2013. [Online]. Available: <http://omf.mytestbed.net/projects/omf6/wiki/Wiki>.
- [7] *Metageek Wi-Spy*, 2013. [Online]. Available: <http://www.metageek.net/products/wi-spy/>.

Testbed Infrastructure for Benchmarking RF-based Indoor Localization Solutions under Controlled Interference

Filip Lemic, Jasper Büsch, Mikolaj Chwalisz, Vlado Handziski and Adam Wolisz
 {lemic,buesch,chwalisz,handziski,wolisz}@tkn.tu-berlin.de
 Telecommunication Network Group (TKN), Technische Universität Berlin (TUB)

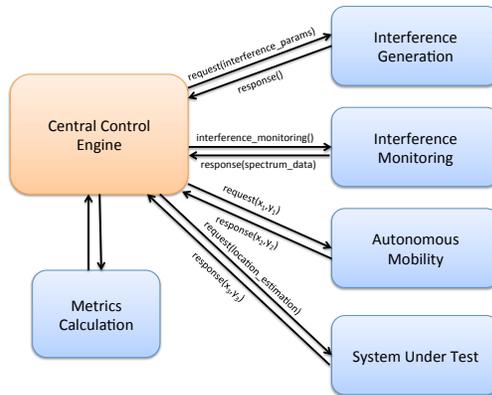
Introduction

Objective benchmarking and comparison of different Radio Frequency (RF)-based indoor localization Systems Under Test (SUTs) is hardly achievable due to:

- Different environments (usually researcher's testbed);
- Different benchmarking scenarios (number and locations of measurement points);
- Unrepeatable conditions in different experiments;
- Unmonitored or marginalized interference effects;

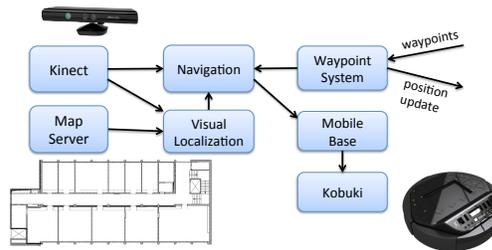
We present an infrastructure for objective benchmarking of RF based indoor localization solutions in an autonomous way without the presence of a test-person.

Testbed Infrastructure



Autonomous Mobility

Navigation is done by correlating given floor map with the distance measurements obtained from the Kinect camera sensor.



System Under Test

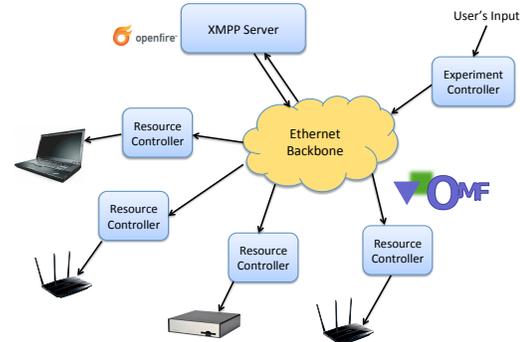
SUT is considered a *black-box* in our approach. The interface that indoor localization SUTs have to provide is rather simple and the only requirement is to report the location estimate when requested. The communication is done using **standard HTTP requests**.

Interference Monitoring

The spectrum in the ISM 2.4 GHz band is free and numerous devices are operating in those frequencies. That is why it is necessary to **monitor if spectral environment looks as expected** in terms of controlled and uncontrolled interference. We can use OMF to orchestrate WiSpy devices, the low cost wireless spectrum analyzers. We mount one device on the robot and one on the interference traffic sink.

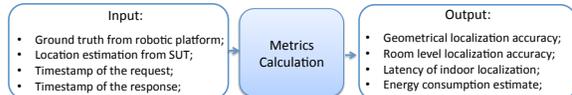
Interference Generation

We have developed means to generate various types of interference scenarios, and in this demo we generate **interference context of typical home or office environments**. Namely, interference is generated using 4 WiFi embedded PCs, representing a server, email client, data client, and video client.



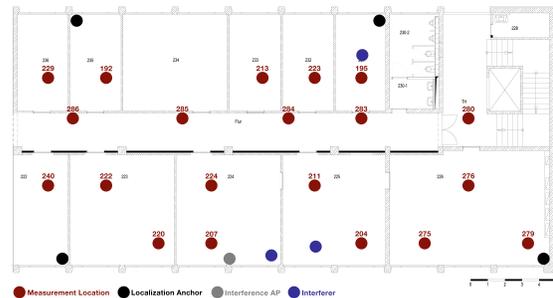
Metrics Calculation

Based on the accurate location obtained from the robotic platform and the location estimation obtained from the SUT, the **engine calculates the performance metrics**:



Demo Description

- Example SUT - RSSI based fingerprinting algorithm;
- Two example experiments:
 - Experiment without controlled interference;
 - Experiment with controlled WiFi interference;



Acknowledgement

This work has been partially funded by the European Commission (FP7-ICT-FIRE) within the projects EVARILOS (grant No. 317989) and CREW (grant No. 258301) and EIT ICT Labs, as part of the activity 12149, "From WSN Testbeds to CPS Testbeds".

EWSN 2014
 The 11th European Conference on Wireless Sensor Networks
 February 17-19, 2014.
 University of Oxford, Oxford, UK

Demo Abstract: Foren6, a RPL/6LoWPAN Diagnosis Tool

Sébastien Dawans, Laurent Deru
CETIC, Belgium
first.last@cetic.be

Abstract—IPv6-based wireless sensor networks are playing an increasingly important role in the development of the Internet of Things. Experience shows that it is not trivial to deploy a large-scale sensor network, despite the many prototyping facilities provided by network simulators, platform emulators and sensor network testbeds. Passive sniffing is a common debugging technique used in real deployments. With Foren6, we propose a novel passive network analyzer aimed specifically at sensor networks operating on emerging IoT standards, 6LoWPAN and RPL, which serves both the academic community and early adopters of 6LoWPAN and RPL.

I. INTRODUCTION

Deploying real-world Wireless Sensor Networks (WSN) has been a challenging issue for many years. Today, a number of tools are available at different phases of a system’s life cycle to reduce deployment risks. WSN simulators and emulators have had a significant impact on application development as well as low-level communication protocol improvements [4]. Likewise, cabled testbeds offer the means to obtain rich debugging information in an environment closer to real-world conditions. Despite such advances, real-world WSNs are still prone to failure.

In a real deployment without the information offered by simulators and testbeds, it is very difficult to debug failures. Some insight on the system conditions can be obtained actively by sending debugging information to a data collection point, but has the drawbacks of draining batteries and being intrusive in the observed systems. For example, actively monitoring a network with link metric estimation problems may remove the problem altogether, while monitoring a dense network may introduce undesirable radio contentions. Additionally, active monitoring will simply fail to report any information in case of network partitioning and other forms of disconnections. An improvement on this is to use a dedicated, high-performance back-channel for propagating network information, but it is also power consuming and requires sophisticated hardware to be interfaced with the deployed nodes [5]. Finally, passive monitoring consists in collecting in-network packet traces to reconstruct the system state without interfering with the network or individual nodes [1].

In this demo, we present Foren6, the only passive monitoring tool to our knowledge which is capable of analyzing RPL-based 6LoWPAN networks. Foren6 uses packet capture devices which we refer to as *sniffers* to reconstruct a topological view of the RPL network in a graphical manner and report

potential network faults and protocol errors. In the remainder of this document, we describe Foren6 and the demonstration details.

II. FOREN6: RPL/6LOWPAN DIAGNOSIS TOOL

Foren6 is a passive monitoring tool for 6LoWPAN networks, published under the GPL licence and freely-available since November 2013. This section covers the most important features of Foren6 related to debugging 6LoWPAN networks in real deployments. Due to space limitations, we refer to the online documentation for illustrations of the user interface [2].

History Navigation 6LoWPAN networks using distance-vector routing protocols such as RPL often go through a network initialization phase, where the routing topology is dynamically formed based on network conditions. When problems arise, it is useful to navigate through the complete network history, including the critical network setup phase, rather than viewing the current state of the network. Foren6 observes control and data-plane traffic to identify individual network states, called *versions*. The entire network state is recorded for each version, allowing to navigate through the network history via a scrollable timeline, while continuing to capture packets in real-time.

Network Visualizer A major contribution of Foren6 is its graphical network representation as shown in Figure 1. The topology offers two complementary modes: data-plane and control-plane. The control plane depicts the RPL DODAGs with arrows representing the current preferred parent selection of each network element. This RPL-related information is enriched with *overlays*, which add color and text to the visualizer according to a certain metric or parameter. Existing overlays include RPL DODAG ID, RPL ranks, total packets sent, DIO/DAO intervals and network faults such as forwarding errors. The data-plane visualizer is a more experimental feature offering insight on application traffic paths.

Acquisition Types As a deployment-oriented debugging tool, Foren6 is designed to capture and process live packets from sniffers. Since it is not always practical to capture and analyze packets at the same time, especially in outdoor deployments, Foren6 can also process packet traces in an offline mode for a post-mortem analysis.

Extended Coverage A passive monitoring tool based on a single sniffer has limited range, even with optimized and costly hardware, and thus cannot capture the entire traffic of large sensor networks. To overcome this limitation, Foren6 proposes

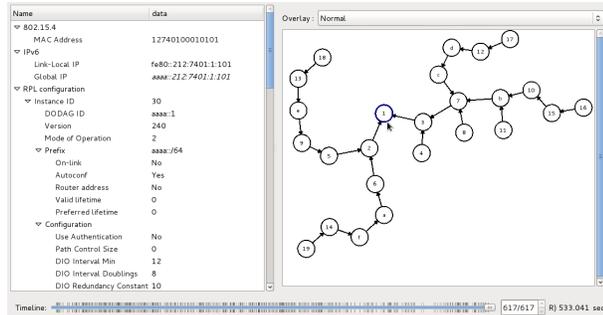


Fig. 1: Foren6 layout in reduced mode with network visualizer and node information. Not shown: event window and packet dissector.

two features: multiple sniffers, and mobile sniffers. First, multiple live sources can be connected to Foren6 simultaneously: Foren6 includes a sniffer synchronization algorithm which uses packets captured by pairs of sniffers. This aggregation algorithm also functions in the offline mode, on traces recorded by multiple sniffers. Second, the tool also supports mobile sniffers since it does not make any assumptions on the sniffers' positions. An Android port of Foren6 is available to analyze a 6LoWPAN network in real-time on a mobile tablet rather than a bulky laptop.

Customizable A minor but highly practical feature is the ability to define a layout to visualize the network topology over the real physical topology rather than using the default node placement algorithm.

III. DEMO DESCRIPTION

In this work, we demonstrate benefits of passive 6LoWPAN networking monitoring in both research and industrial use-cases. From a research perspective, Foren6 has been useful to fix shortcomings in ContikiRPL, the RPL implementation of the Contiki operating system. In more practical situations, this technique helps to pinpoint network deployment problems such as bottlenecks, poorly-connected zones, sensor failures and routing loops.

A. Example 1: Physical network with a Border Router outage

In this first example, we deploy battery-operated Contiki sensors in the demonstration room, and establish two data collection points using two 6LBR-based Border Routers (BRs) in the same RPL Instance [3]. After the initial setup phase, the sensors are spread evenly across DODAGs handled by each BR. A selected sensor continuously pings an IPv6 host on the ethernet backbone, during which we abruptly halt the BR through which the ping traffic is routed. The orphaned nodes will eventually join the remaining DODAG. The network setup, BR outage, transitory forwarding loops and eventual network adaptations will be observed in detail via the different overlays provided by Foren6.

B. Example 2: Offline diagnosis of traffic bottlenecks

Next, the demonstration includes an offline analysis of a previously-recorded packet trace for convenience. This example will observe unequal traffic repartition in a wireless sensor network leading to faster battery depletion than expected on the highly-solicited nodes.

C. Example 3: Diagnosis of ContikiRPL problems in a simulator

This final example shows how passive 6LoWPAN monitoring with Foren6 has recently allowed to pinpoint subtle shortcomings in the multiple-DODAG support of the ContikiRPL implementation, and more generally how Foren6 can be used for improving the communication stack. We deploy a sensor network in a simulated environment using COOJA, which stores packet traces in a file, read by Foren6 on-the-fly via a simple FIFO mechanism. The multi-DODAG issues will be reproduced in COOJA and debugged using Foren6.

IV. CONCLUSION

As a passive 6LoWPAN monitoring tool, Foren6 is useful in both research and industrial deployment settings. In its current state, Foren6 is applicable to RPL and generic 6LoWPAN networks, and has been used mainly in Contiki-based networks. In future work, we plan to monitor other implementations of RPL, leading to more generic decision-making algorithms in Foren6 allowing it to diagnose interoperability issues. There is also a need to extend the tool to 6LoWPAN-ND, allowing to debug networks with a mix of RPL and 6LoWPAN-ND sensors.

ACKNOWLEDGMENT

This work was partly funded by the Walloon Region under the First-DoCA funding number 1017211. We thank Alexis Murzeau and Alexandre Devaux for their valuable contributions.

REFERENCES

- [1] Abdalkarim Awad et al. "On the Need for Passive Monitoring in Sensor Networks". In: *Proceedings of the 2008 11th EUROMICRO Conference, DSD '08*. Washington, DC, USA: IEEE Computer Society, 2008, pp. 693–699.
- [2] S. Dawans and L. Deru. *Foren6 Project Website*. 2013. URL: <http://cetic.github.io/foren6>.
- [3] Laurent Deru et al. "Redundant Border Routers for Mission-Critical 6LoWPAN Networks". In: *Proceedings of the Fifth Workshop on Real-World Wireless Sensor Networks*. 2013.
- [4] Fredrik Österlind. "Improving Low-Power Wireless Protocols with Timing-Accurate Simulation". SICS. PhD thesis. Uppsala University, Department of Information Technology, 2011.
- [5] Matthias Ringwald and Kay Rmer. *Deployment of Sensor Networks: Problems and Passive Inspection*.

Demo Abstract: MPASS: Multi-channel Packet Sniffing System based on IEEE 802.15.4

Jung-Hwan Bae, Tae-Soo Kim, Hiecheol Kim, and Seong-eun Yoo[†]
School of Computer and Communication Engineering
Daegu University, Gyeongsan, Korea
[†]Corresponding author: seyoo@daegu.ac.kr

Abstract—Nowadays there are increasing demands for research on multi-channel based wireless sensor network protocols or applications to support requirements such as increased throughputs or real-time transmissions. The researcher or developer for these protocols or applications have to analyze the exchanged packets in the view points of the correctness of the packet contents and their timeliness at the same time. However, just to use multiple conventional single channel packet sniffers makes the debugging process difficult, for the researchers need to check the correctness of the packets over the multiple channels one by one for themselves. Therefore, we present Multi-channel Packet Sniffing System (MPASS) that helps the multi-channel packet analysis. Wireless packets are detected and timestamped by each sniffer module in MPASS for each channel, and the packets are preprocessed and transmitted to the GUI based analyzer, and then the analyzer parses the received packets and shows them in order. We demonstrate the performance of MPASS by comparing with a widely used packet sniffer.

I. INTRODUCTION

WIRELESS Sensor Networks (WSNs) that build on sensor nodes have been widely spread in military, industrial, medical applications with help of rapidly grown relevant technologies such as semiconductors, embedded software, and so on. Accordingly, the limited ISM band that is commonly used in WSNs becomes too crowded so that the communication protocols based on a simple single channel are difficult to provide some of applications with the sufficient throughputs, reliability, and timeliness. If only a single-channel is used, reliability decreases by the interferences with communication of common or adjacent channels and noise level. Therefore, there are increased research and standardization activities such as IEEE 802.15.4e, ISA100.11a, WirelessHART, and so on to develop the multi-channel communication protocols.

The research and developments for multi-channel communication protocols or applications require convenient and accurate debugging facilities such as a packet sniffer. However, although there are some studies or developments on a single channel packet sniffer, there is lack of study or development of a multi-channel packet sniffer. A way out is to use a set of single channel packet sniffers to analyze a multi-channel protocol. One of severe problems of this approach is the lack of time synchronization between the single channel packet sniffers. Therefore, we propose a Multi-channel Packet Sniffer System (MPASS) which is easily extensible and time

synchronized between sniffer modules. This system can be used to analyze any frames following the standard IEEE 802.15.4 MAC format and IEEE802.15.4 2.4GHz DSSS PHY [1]. We demonstrate the strengths of this system such as scalability and synchronization between sniffer modules by comparing with a well-known packet sniffer.

II. MULTI-CHANNEL PACKET SNIFFING SYSTEM

MPASS consists of a multi-channel packet sniffer (MPS) hardware and GUI (Graphic User Interface) based multi-channel packet analyzer (MPA) software as in Fig. 1. MPS is based on scalable sniffer modules and a time synchronizer. Each of the 16 channels of IEEE 802.15.4 2.4GHz is monitored by each sniffer module, and MPS is easily scaled up to sniff another channel by adding a new sniffer module. The time synchronizer is synchronizing all of the sniffer modules in MPS, and a button in this board gives the stimulus to all of the sniffer modules and makes them synchronized together. The hardware button is used to reduce the uncertainty of the software approach. As long as each sniffer module is synchronized, it captures every packet in its channel, wraps the packet in the UART (Universal Asynchronous Receiver/Transmitter) message format which is explained in Section II-B, and then sends it to MPA. Finally, MPA receives the message from MPS and analyzes it and display the packets in the analysis output window.

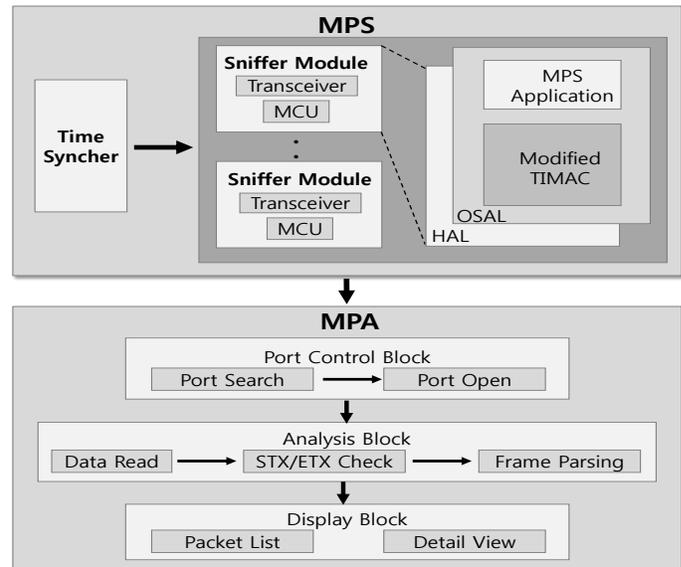


Fig. 1. Multi-channel Packet Sniffing System (MPASS) overview

A. Sniffer Module

Rather than designing a new hardware platform for the sniffer module, we adopt a COTS (Commercial Off-The-Shelf) platform, MSP-EXP430F5438 [2] of Texas Instruments Inc. which consists of a MSP430F5438 MCU and a CC2520EM IEEE 802.15.4 transceiver module [3] as in Fig. 2. One of GPIO (General Purpose Input/Output) pins is taking the signal from the time synchronizer. The software platform of a sniffer module is based on TIMAC 1.4.0, and we modified TIMAC to capture every frame in the radio communication range. To receive the packet regardless of the destination address, MAC radio RX frame filtering function needs to be turned off. Whenever a sniffer module captures each frame, it also takes a timestamp for the frame. The captured frame with the timestamp and the channel number information is passed to the application layer with help of OSAL (OS Abstract Layer) message passing API (Application Programming Interface), and the MPS application encapsulates it by the UART message protocol specified in Section II-B. Finally the sniffer module sends the UART message to MPA via UART.

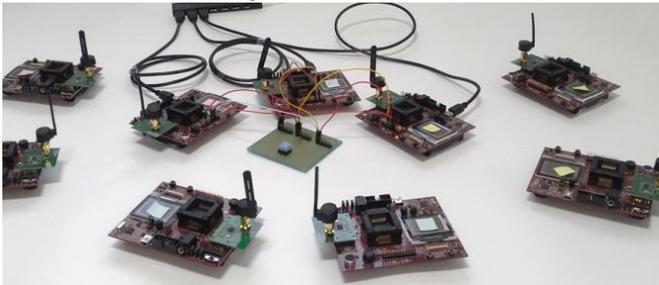


Fig. 2. MPS (3 sniffer modules + time synchronizer) and sensor nodes

B. UART message protocol

When the captured MAC layer frame is sent to MPA, MPS application adds the channel information and the time information which means the synchronized frame reception time. When the MPS application sends the UART message, it needs to manage the start and the end of the message with the special control bytes. 0xAA and 0x7E are used for the start of text (STX) and the end of text (ETX), respectively. In addition, the control escape (ESC), 0x7D is required to discriminate STX, ESC, and STX in the body of a message. If one of these special characters is included in the body of a UART message, ESC is transmitted first and then that the original octet exclusive-or'd with 0x20 is sent. We adopts this framing protocol from RFC1662.

C. GUI based Multi-channel Packet Analyzer

The MPA is designed to receive UART messages from each sniffer module in MPS to analyze the wrapped MAC frames and display them. The MPA consists of port control block, analysis block, and display block as in Fig. 1.

To support the multiple sniffer modules in MPS, MPA is implemented in multi-threaded model. MPA performs in sequence of port scan, port open, and thread run, and the thread repeats to read its own port, analyzes the read message, and displays it in a list control. First, the port control block searches and reads the available port lists in the serial port part of the Windows system registry. Second, the user selects and opens

some of the available ports. When the selected ports open, a thread per each port is created and starts at the same time. If data is ready in an open port, the thread for the port reads the data. When the thread receives the entire UART message, it starts to parse it based on the message protocol specified in Section II-B. The analysis result is shown in the packet list in Fig. 3. The window of MPA is divide into 'Packet List' and 'Detail View'. The packet list is an instance of a list control, and the user can identify each kind of packet types (beacon, data, ACK, and command type) with a unique color for each frame type. It displays the subfields of the message such as channel, time, MAC header (MHR), MAC payload, correlation, link quality Indication (LQI), received signal strength indication (RSSI), and MAC footer (MFR).

'Detail view' shows the detail of a selected packet from the packet list control, and each field of the packet is parsed more in detail to assist the user in analyzing the packet.

Fig. 3. Packet List screen on Multi-channel Packet Analyzer.

III. DEMO SCENARIO

The demonstration consists of a number of sensor nodes using multiple channels, MPASS, and a widely used commercial packet sniffer. MPASS is compared with the commercial packet sniffer in the aspect of the frame capture rate (FCR).

IV. CONCLUSION

To make the research and development on multi-channel communication protocols easier, MPASS architecture and its implementation are described. According to our evaluation, MPASS outperforms a widely used packet sniffer in FCR. Future work includes the research on easier and periodic time synchronization in MPS, the performance improvement in the MPA overloading situation as shown in the evaluation section, and the minimization of the size and the cost of MPASS.

ACKNOWLEDGMENT

This research was supported by Basic Science Research program through the National Research Foundation of Korea (NRF) funded by the Ministry of Science, ICT & Future Planning (No. 2011-0014204) and the IT R&D program of MOTIE/KEIT [10041145, Self-organized Software platform (SoSp) for Welfare Devices].

REFERENCES

- [1] *IEEE Standard for Information technology Telecommunications and information exchange between systems Local and metropolitan area networks Specific requirements*, IEEE Standard 802.15.4, 2006.
- [2] *MSP-EXP430F5438 Experimenter Board User's Guide*, Texas Instruments Inc., DA, Texas, 2013.
- [3] Texas Instruments, *CC2520EMK* [Online]. Available: <http://www.ti.com/tool/cc2520emk>

Demo Abstract: Reliable 6LoWPAN Networks with Multiple Border Routers

Laurent Deru and Sébastien Dawans
 CETIC, Belgium
 first.last@cetic.be

Abstract—Multihop IPv6-based sensor networks support bidirectional internet traffic over tree-like topologies, centered around one or more gateways, or *border routers*. In real-world deployments, there are key benefits to using multiple border routers. This work presents our design and implementation of multiple RPL-based 6LoWPAN border routers in the Contiki operating system and outlines the added benefits in terms of deployment, reliability and energy consumption.

I. INTRODUCTION

The Internet of Things envisions systems comprising large numbers of sensors participating in the IPv6 architecture based on emerging standards and communication protocols. 6LoWPAN and RPL are designed to support a wide range of sensor network applications sharing common constraints of bidirectional multihop communication over energy-constrained nodes [4].

Large battery-operated multihop networks raise challenging deployment issues. Border Routers (BR) are single points of failure, but can be made redundant [2]. Furthermore, far-reaching networks with many hops and a constant traffic per node consume energy unequally across the network: nodes closer to the sink forward more traffic than nodes further down the paths [3].

In this demo, we setup a multi-BR 6LoWPAN/RPL sensor network to demonstrate the added value of having multiple BRs: fault-tolerance, traffic distribution and shorter paths.

II. MULTIPLE BORDER ROUTER DESIGN

Our design of multiple BRs is declined in two approaches. In the first approach, each BR separating the IPv6 and 6LoWPAN networks runs a RPL root to handle the 6LoWPAN side as depicted in Figure 1, top. This can be further divided into a *Router* mode, where two distinct subnets are handled, and a *SmartBridge* mode, where the same IPv6 prefix is shared across either side of the BR. This mode of operation uses active NDP-based synchronization between BRs allowing for seamless interaction with mobile nodes.

In the second approach, an external RPL root runs on the IPv6 network, and each BR is a RPL router (6LR) or a simple switch between Ethernet and 802.15.4 (Figure 1, bottom). No explicit synchronization mechanism is needed in this mode as sensor mobility between BRs is handled intrinsically by RPL.

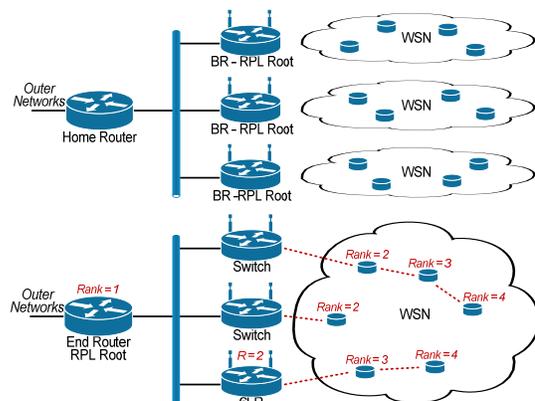


Fig. 1: Top: Routing level aggregation, each BR is a RPL Root with its own DODAG. Bottom: RPL level aggregation, the RPL Root runs on the backbone.

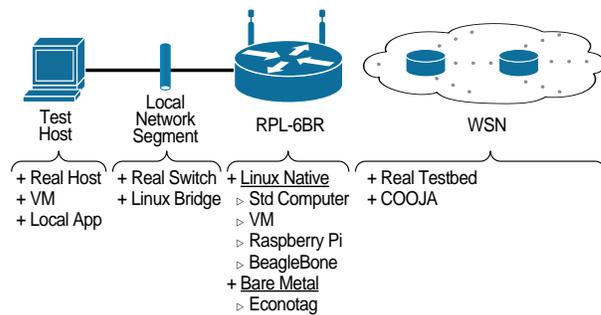


Fig. 2: The modular 6LBR test framework used for the evaluation.

III. DEMO DESCRIPTION

Our demonstration provides insight on how multi-sink sensor networks allow for fault-tolerance, improved traffic distribution and shorter paths. To this end, we conduct three distinct experiments. The first example uses a physical setup to demonstrate how multiple RPL sinks can overcome BR failures. The second example uses a simulated environment to observe how traffic distribution and shorter paths are naturally achieved when deploying multiple BRs over large and widespread networks. The third scenario simulates multi-BRs with mobile nodes.

In both real and simulated WSN scenarios, the BRs are based on our 6LBR implementation in Contiki, compiled for

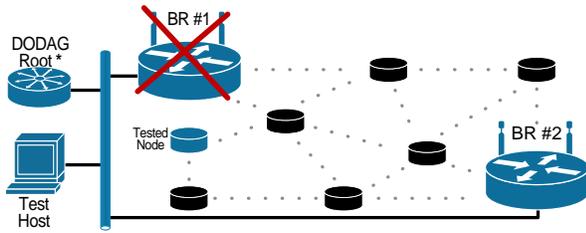


Fig. 3: Border router outage scenario.

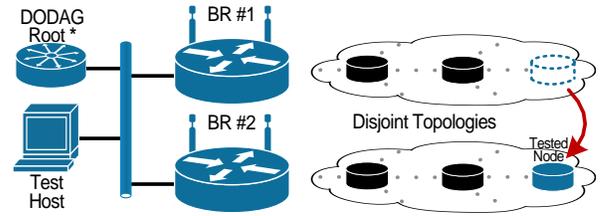


Fig. 4: Mobile sensor scenario.

the Linux platform [1]. For demonstration purposes, we use the modular 6LBR test framework depicted in Figure 2, where individual elements are interchangeable.

A. Example 1: Tolerance to Border Router outages on a real network

The first setup showcases the resilience of the network to BR outages by observing the effect of a BR failure on application traffic routed through the faulty BR. The setup of Figure 2 is used, with a real IPv6 host (Linux PC) connected to two physical 6LBRs running on the BeagleBone platform, through an Ethernet switch. Each 6LBR is placed in the proximity of several battery-operated sensor modules running Contiki with 6LoWPAN and RPL. The experiment is repeated with both routing-level aggregation with the *SmartBridge* mode and RPL-level aggregation with an external RPL root.

After an initial network setup phase, one of the sensor modules initiates periodic Ping messages destined to the IPv6 host (the test PC). Once the network is stable, the BR on the path of this Ping traffic is abruptly halted, causing a part of the network to lose all connectivity to the outside network as shown in Figure 3. Orphaned nodes eventually attach themselves to the nodes connected to the remaining BR using ContikiRPL’s multiple-DODAG support for the *SmartBridge* mode, and basic intra-DODAG functionalities in the case of an external RPL root. The demo shows that the downtime incurred by the BR outage lasts only a few seconds in the presence of periodic application traffic.

B. Example 2: Performance gains in a large, simulated network

The second example leverages the COOJA simulator to scale to large networks. In this scenario, the 6LBRs run as Linux processes on the test machine, and are connected to each other and the host via a virtual bridge in Linux. The 6LBRs are interfaced with the sensors in COOJA via a simulated 802.15.4 interface.

Two almost identical simulations are executed in parallel: the first one consists of many sensor nodes spread out in COOJA interfaced with a single BR and reporting periodic UDP messages to an external server. The second simulation implements the exact same scenario, with an extra BR at the opposite side of the sensor network. In each simulation, a graphical window displays important statistics from an application and infrastructure perspective: average packet latency and per-node energy consumption.

The dual-BR simulation has a lower latency thanks to a lower average hop-count. Its energy consumption is more evenly spread among network nodes due to the greater path diversity offered by multiple sinks.

C. Example 3: Inter-BR node mobility

The last example demonstrates how our design supports inter-WSN node mobility. Two BRs connected on a common Ethernet backbone are interfaced two disjoint simulated sensor networks. A mobile node periodically moves from one WSN to the other while exchanging a constant Ping traffic with an external host as in Figure 4. The demo measures the downtime each time the mobile node moves out of radio range of its current DODAG.

IV. CONCLUSION

The architecture we present in this work addresses important practical deployment issues by leveraging multiple BRs. We show that supporting multiple BRs is simple and does not necessarily rely on sophisticated synchronization mechanisms. This multiple BR design improves overall reliability, supports mobile nodes and results in measurable energy savings in large networks.

ACKNOWLEDGMENT

This work was partly funded by the Walloon Region under the First-DoCA funding number 1017211.

REFERENCES

- [1] S. Dawans, M. Ocaña, and L. Deru. *RPL-6LBR implementation in Contiki* <http://cetic.github.com/6lbr>. tag 6lbr-1.2. 2013.
- [2] Laurent Deru et al. “Redundant Border Routers for Mission-Critical 6LoWPAN Networks”. In: *Proceedings of the Fifth Workshop on Real-World Wireless Sensor Networks*. 2013.
- [3] Leila Ben Saad and Bernard Tourancheau. “Multiple Mobile Sinks Positioning in Wireless Sensor Networks for Buildings”. In: *Proceedings of the Third International Conference on Sensor Technologies and Applications*. SENSORCOMM ’09. IEEE Computer Society, 2009.
- [4] T. Winter (Ed.), P. Thubert (Ed.), and RPL Author Team. *RPL: IPv6 Routing Protocol for Low power and Lossy Networks*. Mar. 2012.

POSTER ABSTRACT: DECENTRALIZED TIME-SYNCHRONIZED CHANNEL SWAPPING FOR WIRELESS SENSOR NETWORKS

G. Smart, N. Deligiannis, Y. Andreopoulos

University College London, UK
Dept. of Electronic Engineering

R. Surace, V. Loscri, G. Fortino

University of Calabria, Italy and INRIA, France
Dept. Electr., Info. and Syst. (IT) and FUN Team (FR)

ABSTRACT

We are working on a new concept for decentralized medium access control (MAC) coordination, termed as *decentralized time-synchronized channel swapping* (DT-SCS). Under the proposed DT-SCS and its associated MAC-layer protocol, wireless nodes converge to synchronous beacon packet transmissions across all IEEE802.15.4 channels with balanced number of nodes in each channel. This is achieved by the proposed reactive listening mechanisms, based on pulse coupled oscillator techniques. Peer-to-peer *channel swapping* can then take place via swap requests and acknowledgments made by concurrent transmitters in neighboring channels. Comparisons of DT-SCS against time-synchronized channel hopping (TSCH) reveal that our proposal comprises an excellent candidate for completely decentralized MAC-layer coordination in WSNs by providing for quick convergence to steady state, high bandwidth utilization, high connectivity and robustness to interference and hidden nodes.

1. INTRODUCTION

The concept of channel hopping has gained acceptance as a good solution according to high-bandwidth, energy-efficient, wireless sensor networks (WSNs), with time synchronized channel hopping (TSCH) [1] now being part of the IEEE802.15.4e-2012 standard. Via the advertising mechanism of TSCH, each node reserves timeslots within the slotframe interval and within the 16 channels of IEEE802.15.4. As the slotframe interval repeats periodically, all nodes transmit and listen in different channels, thus avoiding concentrated interference. However, the TSCH slotframe has a rigid (pre-defined) structure and filling up the available slots follows a rather complex advertising request and acknowledgment (RQ/ACK) process on a coordination channel. This channel is prone to interference and occasional self-inflicted collisions when the nodes are set to advertise slot reservations very aggressively. Conversely, if slot advertising is not aggressive and nodes leave the network, their slots may remain unoccupied for long periods until another advertisement RQ/ACK process reassigns them to other nodes. This limits the bandwidth usage per channel.

Our work addresses these issues based on the concept of *pulse coupled oscillators* (PCOs) [2, 3]. Specifically, we propose a novel *decentralized time-synchronized channel swapping* (DT-SCS) framework, in which nodes randomly join a channel and achieve PCO-based coordination via the periodic transmission of *beacon* packets at the MAC layer. For channels with equal number of nodes, DT-SCS converges to synchronized beacon packet transmission at the MAC in a completely uncoordinated manner. Furthermore, it allows for arbitrary pairwise swaps between nodes in neighboring channels with limited effort and without disrupting the WSN operation. Finally, due to the inherent adaptation of PCO mechanisms to the effects of nodes joining and leaving the process, our proposal is robust to interference as well as node churn during WSN reconfiguration.

2. PCO-BASED SYNC/DESYNC FOR DT-SCS

Consider a WSN consisting of W nodes randomly distributed in C channels, with each node transmitting short *beacon* packets periodically every T seconds. In the proposed DT-SCS mechanism, the nodes in each channel perform PCO-based desynchronization (i.e., they are “DESYNC” nodes) and elect a single “SYNC” node to provide for cross-channel synchronization. Within each period, the SYNC node of each channel listens for the SYNC beacon message in the next channel (with cyclic behavior between channels 1 and 16) and adjusts the transmission time of its own beacon packet in its own channel using PCO-based synchronization [3]. SYNC nodes will also move to the next channel if they detect that less nodes are present there. In this way, the WSN can converge to the steady state with $W_c = \frac{W}{C}$ nodes per channel (when W is not divisible by C , the case the scheme balances to $W_c \in \{\lfloor \frac{W}{C} \rfloor, \lceil \frac{W}{C} \rceil\}$ nodes per channel). The beacon packet transmission flow between DT-SCS nodes is schematically illustrated in Fig. 2 for the case of $C = 3$ channels with $W_c = 4$ nodes per channel (i.e., $W = 12$).

Once the system reaches the steady state, SYNC or DESYNC nodes in adjacent channels can swap channels and timeslots in pairs using a simple RQ/ACK scheme within a short predefined guard time before and after the expected

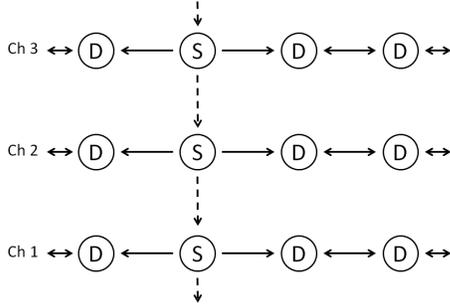


Fig. 1. DT-SCS within 3 channels, showing the intra-channel desynchronization (solid lines) and inter-channel synchronization (dotted lines) between DESYNC (D) and SYNC (S) nodes, respectively. Arrows indicate the intended recipient of each beacon packet transmission.

beacon transmission. If nodes join or leave the network, all remaining nodes adjust their beacon packet timings spontaneously, in order to converge to a new steady state.

Once convergence to steady state is achieved, the only overhead in the proposed DT-SCS protocol stems from handling swap requests as well as beacon packet broadcasts. Both, however, are very short packets (less than ten bytes), which makes the overhead minimal compared to the payload packet transmission and reception.

The loss of beacon packets and timing errors due to interference cause node beacon times to *waver*, i.e., nodes send beacon messages at incorrect times. As such, all nodes receiving these messages are similarly affected. To combat this, we consider the notion of *coupling* between nodes, introduced by PCOs [3]: instead of a DESYNC node jumping directly to the midpoint of its beacon neighbors, the node *slides* towards the mid point with coupling factor α ($0 < \alpha < 1$); similarly, a SYNC node gradually adjusts its beaoning time by coupling factor β ($0 < \beta < 1$) to align with the beacon of the SYNC node in the next channel. Using PCOs with appropriate coupling factors ensures that any noise and instability in beacon timings is attenuated and does not propagate uncontrollably throughout all nodes and channels of DT-SCS.

3. INDICATIVE EXPERIMENTS WITH TELOS B MOTES

We implemented DT-SCS as an application in the Contiki 2.6 operating system running on TelosB motes. By utilizing the NullMAC and NullRDC netstack options of Contiki, we control all node interactions at the MAC layer via our application code. We consider a WSN deployment with 64 nodes in 16 channels, which leads to $W_c = 4$ nodes per channel in the steady state. We set the beacon period to $T = 228$ ms.

To measure the energy consumption of DT-SCS per node, we placed a TelosB sensor running DT-SCS in series with a

high-tolerance 1-Ohm resistor and utilized a high-frequency oscilloscope to capture the current flow through the resistor in real time. Average results collected over 10 minutes of operation are reported. The average power consumption without transmitting or receiving payload was measured to be 13.57 mW. This is mainly due to maintaining beacon transmission and reception. Importantly, once convergence is achieved, the power overhead can be reduced by listening for beacon packets less frequently. For instance, setting the nodes to listen for beacons every eight periods brings the power consumption down to 1.58 mW. To set an illustrative comparison, the power consumption of a node operating TSCH under minimal payload (i.e., 128 bytes per four seconds) is 1.64 mW [4].

Under the aforementioned conditions, the convergence of the WSN using the proposed DT-SCS protocol was achieved within 3.73 s on average. On the other hand, TSCH convergences within 14.00 s on average (under the 6tisch simulator [4]). In addition, we measured the maximum achievable network throughput (i.e. transmission rate by all nodes in the WSN when they utilize their entire slot) to be 94.02 kbps. This is significantly higher than the corresponding maximum achievable network throughput of TSCH, which, under the default 6tisch setup [4], was measured to be 57.50 kbps.

4. CONCLUSION

The unique aspect of our approach is the concurrent use of pulsed coupled oscillators that perform synchronization and desynchronization in multiple channels. This allows for rapid convergence to the steady state in a completely decentralized manner, that is, without requiring a node or channel coordinator, or time synchronization via a global clock. Experimentation via simulations and a real implementation shows that, in comparison to TSCH, the proposed DT-SCS leads to a significant reduction of the convergence time and substantially higher network utilization.

5. REFERENCES

- [1] T. Watteyne et al., “OpenWSN: a standards-based low-power wireless development environment,” *Trans. on Emerg. Telecom. Techn.*, vol. 23, no. 5, pp. 480–493, 2012.
- [2] J. Degeysys et al., “Desync: self-organizing desynchronization and tdma on wireless sensor networks,” in *Proc. IPSN. ACM*, 2007, pp. 11–20.
- [3] R. Pagliari, Y.-W. P. Hong, and A. Scaglione, “Bio-inspired algorithms for decentralized round-robin and proportional fair scheduling,” *IEEE J. on Select. Areas in Commun.*, vol. 28, no. 4, pp. 564–575, May 2010.
- [4] X. Vilajosana et al., “A realistic energy consumption model for tsch networks,” *IEEE Sensor J.*, 2013.

Poster Abstract: Fairness in WSNs

Yad Tahir and Julie A. McCann

Department of Computing, Imperial College, London, UK

Email: {yst11, j.mccann}@imperial.ac.uk

Abstract—Wireless sensor networks (WSNs) are offering tremendous benefit for a large range of systems and industries. However, WSN remain costly in terms of provisioning and maintenance. To alleviate these costs we envisage that such infrastructures will be shared among concerns; a notion which is supported by the types of hardware being deployed in the real world having increased resources to do so. In such systems nodes must therefore meet the requirements of multiple users; users require different ratios of the different resources. This poster looks at the problem of multi-resource fairness in WSNs and the particular challenges that WSN present in this context.

I. INTRODUCTION

There is currently a large movement to make future cities sustainable. This sustainability is driven from the efficiencies that can be made in city processes which in turn are informed by the knowledge obtained from sensor systems monitoring city infrastructures. Therefore to achieve this many sensor networks will be deployed, but if this happens organically, a given space will be covered by many WSN unnecessarily. This is especially wasteful given that the deployment and maintenance of a WSN is still very costly. Further, this has implications in terms of the maximum capacity that such networks can achieve in this dense bounded space as wireless network collisions will increase proportionally. In parallel we are seeing an increase in the systems capabilities that compose WSN, i.e the nodes being deployed in the real world are not like those predicted over 10 years ago. Putting these factors together provides a strong argument for WSN sharing whereby different concerns share the multi-purposed sensor network. In fact it has been already found that multi-purpose WSNs serving more than one application are remarkably more cost efficient than sensor networks dedicated to a single (global) application [1]. The two core challenges here are concerned with resource sharing and security. In this paper we focus on the former.

Deploying more than one application on a WSN can cause hardware inconsistency and resource starvation. Even with advances in embedded hardware capacities, nodes in WSNs still have a very limited amount of resources. Badly designed, or malware, applications can consume the available resources ultimately causing resource starvation for other applications. While using modern embedded operating systems such as Contiki and TinyOS can guarantee the consistency of peripheral devices via locks (i.e. not allowing a process to use the device when it is busy), ensuring *fair* resource utilization amongst applications is yet to be addressed in WSNs.

The notion of *fairness* is not new. In fact, it has been studied widely in economics. Although fairness appears very simple

to understand, applying it in computer systems can be difficult as the resource allocation mechanism can be directly related to the subjective preferences of users.

One of the approaches to obtain resource fairness is via task scheduling; a mechanism that has been around for some time to share resources in multitasking and multiprogramming operating systems. However, traditional *preemptive* scheduling algorithms such as the Completely Fair Scheduler (CFS) can only achieve *single-resource fairness* (i.e. it ensures that each process gets a fair share of one resource type - CPU). Achieving *multi-resource fairness* can be through a resource allocation mechanism that sits on top of the operating system layer. Here, the allocation policy mechanism performs process multiplexing/forking operations to assign resources to users. However, such approaches are not appropriate for WSN operating systems.

To achieve this we are investigating the Dominant Resource Fairness (DRF) allocation scheme [2] which is used in more traditional computing systems where multiple types of resources are required to be shared in a fair way. Unlike other resource allocation solutions, DRF has several desirable properties, including *strategy-proofness*, *sharing-incentive*, *envy-freeness* and *pareto-efficiency*. These properties are beginning to be proven effective in traditional (Data Centre scale) computing [2] and Queuing devices (e.g. Middleboxes) [3]. However, to the best of our knowledge, there is no work applying DRF to WSNs or other networked embedded systems as yet.

Through initial experimentation, we show that it is not a simple task adapting the current work on DRF allocation to WSN for a number of reasons discussed in the next section.

II. FAIR RESOURCE USAGE USING SCHEDULING

Many embedded operating systems in WSNs use *non-preemptive* scheduling algorithms. For instance, Contiki uses the FCFS scheduling policy to schedule runnable protothreads[4]. This means that the Contiki scheduler cannot ensure fair MCU usage between the protothreads and long CPU-burst processes can cause MCU resource starvation. Any notions of fairness would have to be directly coded in the current running process as only it has the ability to release the resources it owns at that time.

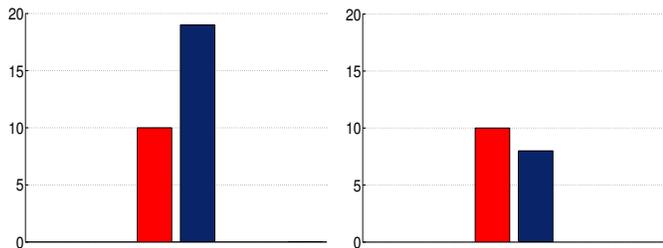
Utilizing a traditional or more appropriately lightweight preemptive scheduling policy such as Round-robin (RR) can certainly solve this problem by assigning equal time slices for each protothread. By doing so, RR guarantees that each protothread gets an equal amount of MCU share, but the MCU

is just one resource. Applications in wireless sensor networks use more than one resource type such as wireless bandwidth and energy and each application may have different demands for each of these resource types.

DRF cannot be implemented directly in wireless sensor networks. The default DRF implementation is not suitable for memory-limited devices as it uses process multiplexing techniques to allocate resources. Most modern embedded operating systems do not support process forking. More importantly, adding another resource management layer is costly.

Therefore, we propose our Dominant Resource Scheduling (DRS) scheme as a solution to this problem. DRS a resource-aware scheduler designed to achieve dominant resource fairness by dynamically adjusting the MCU share to applications. Instead of performing multiplexing operations, in DRS the fair usage of multiple resources is ensured during task scheduling.

Thus, unlike other scheduling solutions, DRS achieves *multi-resource fairness* and has the attractive (yet important) properties of DRF including *strategy-proofness*, preventing users obtaining benefit from lying or adjusting their true demand vectors. In fact, based on our initial real sensor node experiments (written for Contiki OS) we often notice users lying about their resource requirements causes to have worse resource allocation as shown in Fig. 1. In the experiments, a Tmote Sky node is attached to a 3G module (SIM5218) allowing users to send 200 sms every hour. The node also has a solar panel with 150 mWh energy budget[5]. Lying about resource demand ratios causes the user to have a lower MCU share which proportionally reduces the number of tasks the user can execute in every duty cycle.



(a) Before lying: The demand vectors of user1 and user2 are (7mW, 10sms), (4mW, 1sms)
 (b) After lying: The demand vectors of user1 and user2 are (7mW, 10sms), (10mW, 1sms)

Fig. 1: Number of tasks executed by user1 (Red) and user2 (Blue) before and after lying.

Our initial implementation show that there are potentially large benefits using DRF with DFS and therefore for the subsequent parts of this PhD research we wish to answer the following questions:

- There is a fundamental trade-offs between system fairness and efficiency (e.g. throughput) [6]. Recently, [7] has shown that DRF has poor efficiency performance (in terms of social welfare). The utility function proposed in [6] aims to balance the efficiency-fairness trade-offs, but it is too heavy for resource-limited WSNs. *How can*

we implement a new lightweight approach to balance the trade-off between system efficiency and fairness for multi-purposed WSNs?

- Correlations between resources are ignored by the current state-of-the-art DRF based allocation policies (even where DRF is being used in traditional computing systems). In WSNs, resources are very correlated. For instance, there is a natural relationship between energy consumption and bandwidth usage. *How can we provide fair scheduling which not only ensures fairness of multiple resource types, but also considers the correlations that exist between the resources?*
- Some sensing applications require delay guarantees. Single-resource scheduling accounting for task deadlines (i.e. real-time scheduling) has been well-studied. However, temporal guarantees have not been studied for DRF in traditional computing yet alone WSN. *Can we ensure fairness with soft deadlines through partitioning tasks and rescheduling sub-tasks? In addition, how can we design a unifying framework addressing the trade-off between multi-recourse fairness and delay guarantees?*

III. CONCLUSION & FUTURE WORK

This poster investigates the fairness problem that exists in networked embedded systems such as WSNs. We briefly explain why this problem should be considered in multi-user networks. We clearly state key research questions that must be addressed in this topic area. We also explain how scheduling can be used to achieve single-resource or multi-resource fairness in WSNs.

We present a first study of DRF in the context of WSNs. We explain why DRF cannot be ported directly on top of most common WSN operating systems. We propose a novel scheduling algorithm, Dominant Resource Scheduling (DRS), which adaptively adjusts the MCU share to ensure fair usage of other resource types. Although DRS successfully implements the idea of DRF in the context of WSNs, many problems remain, including spatial fairness, correlation-aware solutions, and deadline requirements. These issues will be addressed in our future work.

REFERENCES

- [1] A. H. Akbar, A. A. Iqbal, and K.-H. Kim, "Binding multiple applications on wireless sensor networks," in *GPC*, 2006, pp. 250–258.
- [2] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica, "Dominant resource fairness: fair allocation of multiple resource types," in *USENIX NSDI*, 2011.
- [3] A. Ghodsi, V. Sekar, M. Zaharia, and I. Stoica, "Multi-resource fair queueing for packet processing," *ACM SIGCOMM*, pp. 1–12, 2012.
- [4] A. Dunkels, O. Schmidt, T. Voigt, and M. Ali, "Protothreads: simplifying event-driven programming of memory-constrained embedded systems," in *Proc. ACM SenSys*, 2006, pp. 29–42.
- [5] S. Yang, X. Yang, J. A. McCann, T. Zhang, G. Liu, and Z. Liu, "Autonomic solar powered wsn for network-wide protocols," in *IEEE JSAC*, to appear.
- [6] C. Joe-Wong, S. Sen, T. Lan, and M. Chiang, "Multi-resource allocation: Fairness-efficiency tradeoffs in a unifying framework," in *Proc. IEEE Infocom*, 2012, pp. 1206–1214.
- [7] D. C. Parkes, A. D. Procaccia, and N. Shah, "Beyond dominant resource fairness: extensions, limitations, and indivisibilities," in *Proc. ACM Conference on Electronic Commerce*, 2012, pp. 808–825.

Poster Abstract: AGRO - Optimal Routing with Low Latency and Energy Consumption in Wireless Sensor Networks with Aerial Relay Nodes

Jin Zhang, Kai Li, Andreas Reinhardt, Salil S. Kanhere
 School of Computer Science and Engineering
 The University of New South Wales, Sydney, Australia
 Email: {jinzhang, kail, andreas, salilk}@cse.unsw.edu.au

Abstract—Energy efficiency is always a challenge in Wireless Sensor Networks (WSNs). Unmanned Aerial Vehicles (UAVs) can relay the data to a remote base station, which can help save energy for the ground sensors in the WSN. However, the aerial link in a UAV network is not always connected due to the movement of UAVs. As a result, routing data in the UAV network causes significant delay. To achieve maximum throughput with lowest latency and energy consumption, we propose an Air-Ground Routing Optimization (AGRO) which selects the optimum routing paths from all possible Air-Air, Air-Ground, Ground-Air or Ground-Ground links at each hop.

I. INTRODUCTION

Wireless sensors are often deployed in a large area or in harsh environments for disaster recovery, environmental monitoring or target-tracking. In these scenarios, sensors transmit their data to the base station in multiple hops. In order to achieve energy efficiency in WSNs, UAVs equipped with flight control systems and radio transmitters have been used to relay data for the ground sensor nodes [1], [2]. UAVs significantly prolong the lifetime of the WSN since they have large radio coverage and their battery can be recharged conveniently. A UAV collects the data from the ground sensors and acts as a relay node in a multi-hop aerial network as shown in Figure 1. We assume the WSN on the ground is connected. Hence, the routing latency in the WSN is negligible compared to the UAV network.

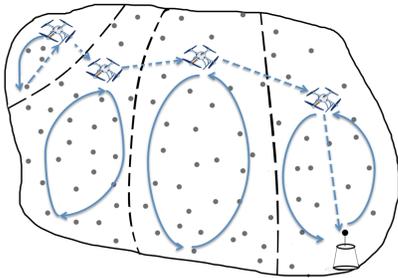


Figure 1: The structure of a UAV-WSN system. Each UAV covers a part of ground area.

However, incorporating UAVs in WSNs presents several non-trivial challenges on data routing. Usually, the UAV network is not always connected in the UAV-WSN system. Since the transmission power of ground sensors is low, UAVs need to move in different areas separately so as to cover

all ground sensors. Therefore, if a ground sensor node only chooses the UAV as the next hop but it is not in the radio coverage of the UAV, it has to buffer its data and wait for the UAV for a long time to relay the data to the base station.

To address these issues, in this preliminary work, we propose AGRO which finds a route to the base station by selecting either Air-Air, Air-Ground, Ground-Air or Ground-Ground link at each hop so that more data is received with low latency and energy consumption of the ground sensor nodes, considering the network that includes both static ground sensors and mobile UAVs.

II. PROBLEM STATEMENT AND AGRO

The total number of ground sensor nodes is n and we have m UAVs in the air ($n > m$). Given a combination of Air-Air, Air-Ground, Ground-Air, Ground-Ground links, the routing problem in the UAV-WSN system is how to find an optimal path so that the data reception at the base station is maximized during a data collection time while the energy consumption and routing latency are both upper bounded. We assume the ground sensor nodes are the only data sources and UAVs only relay data. UAVs move at constant speed following predefined trajectories. Specifically, the connectivity between two UAVs depends on their trajectories, which means the link is connected only when one UAV flies into the radio range of the other. We define the initial energy of a ground sensor node i as E_i^0 and the data generation rate of the ground sensor node i as λ_i . In order to prevent a ground sensor node from completely depleting its battery, we assume that it powers down if its residual energy goes below a certain threshold $E_{td} > 0$. The energy consumption of the ground sensor node i on routing path \mathcal{P} is defined as $\check{E}_i(\mathcal{P})$. We assume both UAV and base station have unlimited energy. $\check{E}_i(\mathcal{P})$ contains the energy consumption on transmitting packets to the neighbor i'' and receiving packets from its neighbor i' . Hence,

$$\check{E}_i(\mathcal{P}) = (P_{(receiving)ii''} + P_{(sending)ii'}) \cdot T_i(\mathcal{P}) \quad (1)$$

$T_i(\mathcal{P})$ is the time for the ground sensor node i to transmit and receive data packets on the routing path \mathcal{P} . We define the number of ground sensor nodes on the routing path in

the WSN is $n(\mathcal{P}_g)$ and the energy consumption of the path \mathcal{P}_g is $\sum_{i=1}^{n(\mathcal{P}_g)} \check{E}_i(\mathcal{P}_g)$.

We define the waiting time of neighbor node j' for UAV j in aerial routing path \mathcal{P}_a to cover its own area as $\tau_{jj'}(\mathcal{P}_a)$. As a result, the maximum waiting time for ground sensor node i is $\sum_{j=1}^{m(\mathcal{P}_a)} \tau_{jj'}(\mathcal{P}_a)$.

Objective Maximize $\sum_{i=1}^{n(\mathcal{P})} (\lambda_i \cdot T_i(\mathcal{P}))$

Constraints

- $\sum_{i=1}^{n(\mathcal{P})} \check{E}_i(\mathcal{P}) \leq \sum_{i=1}^{n(\mathcal{P}_g)} \check{E}_i(\mathcal{P}_g)$
- $E_i^0 - \sum_{i=1}^{n(\mathcal{P})} \check{E}_i(\mathcal{P}) \geq E_{td}$
- $\sum_{k=1}^{n(\mathcal{P})} \sum_{i=1}^k \tau_{ik}(\mathcal{P}) \leq \sum_{j=1}^{m(\mathcal{P}_a)} \tau_{jj'}(\mathcal{P}_a)$

The proposed optimization is modeled as a linear programming problem and can be solved in linear time. In AGRO, the optimal routing path \mathcal{P} is calculated at the base station. Since the flight trajectory of UAV is fixed and known, \mathcal{P} can be pre-calculated. The routing path \mathcal{P} indicates the next-hop link in the routing table locally stored at each node. AGRO makes use of routing tables to select the next-hop link so that the packet can be transmitted.

III. PERFORMANCE ESTIMATION

Based on the setting in a real-world application [3] where the authors use the UAV to monitor the bio-geophysical properties of the vines in a vineyard, we define the following parameters for the energy consumption and latency estimation:

- 11 ground sensor nodes deployed in a 200m x 300m area shown in Figure 2. The area is divided to 3 subareas equally so that 3 UAVs can cover the whole area shown in Figure 3.
- Velocity of UAV is $2.4m/s$ and one round flight time of a UAV to cover its area is $T_u = 5mins$.
- The size of one data packet is 32 bytes and total data payload of one ground sensor node is 1MB.

Firstly, we estimate the upper bound of energy consumption. If the source node sends data to the base station and there are 11 hops on the ground, the energy consumption of 10 relay ground sensor nodes on the routing path is $\check{e} * \frac{1024^2}{32} * 10 = 49.152J$ [4]. \check{e} is the energy consumption on receiving and transmitting one data packet. Secondly, we estimate the upper delay bound as follows. In the worst case, a ground sensor node at the left bottom corner of the subarea A generates data when *UAV-1* just flies past. In this case, the ground sensor node needs to wait $2 * T_u$ to route data to *UAV-1*. The total delay is $2 * T_u + T_u + T_u = 20mins$. Based on the estimated energy consumption and time delay, there exists an optimal routing path that contains both air and ground links to achieve lower latency and more energy efficiency. For example, in Figure 2, *UAV-1* and *UAV-2* save the energy for the ground sensor node *S3* since they relay the data from *S2*. In Figure 3, UAVs cause the latency on the data routing if they are selected as next hop.

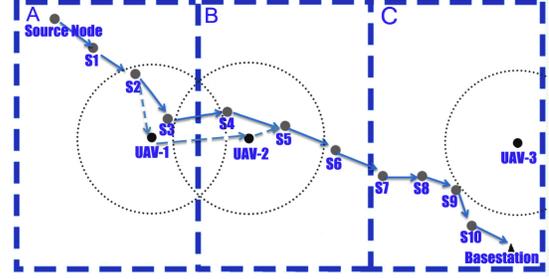


Figure 2: Estimating the upper bound of energy consumption of the ground WSN

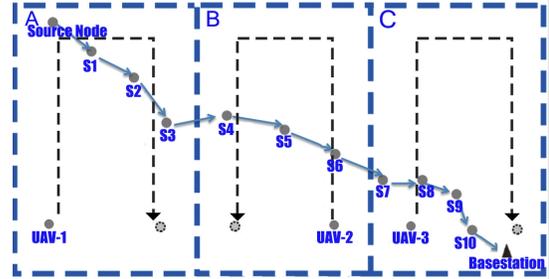


Figure 3: Estimating the upper bound of latency introduced by UAVs

IV. CONCLUSION

The work presented in this paper focuses on the energy efficiency and routing latency problem in a WSN-UAV system. Previously UAVs were assumed to cause much routing delay to WSNs. We have thus presented AGRO that significantly reduces the energy consumption of ground sensor nodes while guaranteeing low latency, by flexibly selecting a combination of Air-Air, Air-Ground, Ground-Air or Ground-Ground link at each hop. We will implement the linear solution as well as heuristics for the analysis of data reception, energy efficiency and latency in our future work.

REFERENCES

- [1] D.-T. Ho and S. Shimamoto, "Highly reliable communication protocol for wsn-uav system employing tdma and pfs scheme," in *GLOBECOM Workshops*. IEEE, 2011, pp. 1320–1324.
- [2] J. A. Cobano, J. Martínez-de Dios, R. Conde, J. Sánchez-Matamoros, and A. Ollero, "Data retrieving from heterogeneous wireless sensor network nodes using uavs," vol. 60, no. 1. Springer, 2010, pp. 133–151.
- [3] D. Turner, A. Lucieer, and C. Watson, "Development of an unmanned aerial vehicle (uav) for hyper resolution vineyard mapping based on visible, multispectral, and thermal imagery," in *Proceedings of 34th International Symposium on Remote Sensing of Environment*, 2011, p. 4.
- [4] K. Li, B. Kusy, R. Jurdak, A. Ignjatovic, S. S. Kanhere, and S. Jha, "k-fsom: Fair link scheduling optimization for energy-aware data collection in mobile sensor networks," The University of New South Wales, Australia, Tech. Rep., 2013.

Poster Abstract: Neighborhood Cardinality Estimation in Extreme Wireless Sensor Networks

Marco Cattani* Marco Zuniga Andreas Loukas Koen Langendoen

{m.cattani, m.a.zunigazamalloa, a.loukas, k.g.langendoen}@tudelft.nl
Delft University of Technology, The Netherlands

Abstract—We address the problem of estimating the neighborhood cardinality of nodes in extreme wireless sensor networks, where nodes are mobile, densities reach hundreds of neighbors, and all nodes want to estimate cardinality concurrently. Estimators need not only be accurate, but also fast, asynchronous (due to mobility), and lightweight (due to concurrency and high density). To cope with these requirements, we propose a mechanism with extremely low overhead that leverages the rendezvous time of duty-cycled medium access control (MAC) protocols; the shorter the rendezvous, the higher the cardinality, and vice versa. Implemented on Contiki OS and deployed on a 100-node testbed, our estimator determines neighborhood cardinalities with less than 10% error.

I. MOTIVATION

Our work is part of a larger project (COMMIT/) related to public safety. The goal is to provide coin-sized devices to attendees in large-scale open-air festivals, and issue alerts when the crowd *density* crosses a dangerous threshold. In these scenarios, all devices need to periodically estimate their surrounding density, which drastically fluctuates, reaching peaks of hundreds of nodes. State-of-the-art estimators achieve an accuracy between 3% and 35% for up to 25 smartphones, relying on audio signals [1], bluetooth signals [2], or radio signals [4]. However, only a fraction of the nodes perform the estimation process, rendering them unsuitable for our application. In this work we therefore propose a low-overhead cardinality estimator that is robust to high densities and mobility, and show that it scales from few tens to one hundred neighbors, while allowing all nodes to perform the estimation concurrently.

II. MECHANISM

The key idea of our estimator is simple: in duty-cycled MAC protocols nodes periodically wake up within a given period (t_w) to listen for incoming messages. The time difference between two consecutive nodes' wake-ups (*rendezvous time*) captures the density of the neighborhood (n). The shorter the rendezvous time, the higher the density, and vice-versa. According to order statistics, the rendezvous time can be modeled using a Beta random variable U_k with expected value equal to $t_w/(n+1)$. Inverting this expectation, we can use the average of the observed rendezvous \bar{u}_k to estimate the neighborhood cardinality $\hat{n} = (t_w/\bar{u}_k) - 1$.

III. IMPLEMENTATION

We implemented our estimator using a slightly modified version of Low Power Listening (LPL). As usual, when a

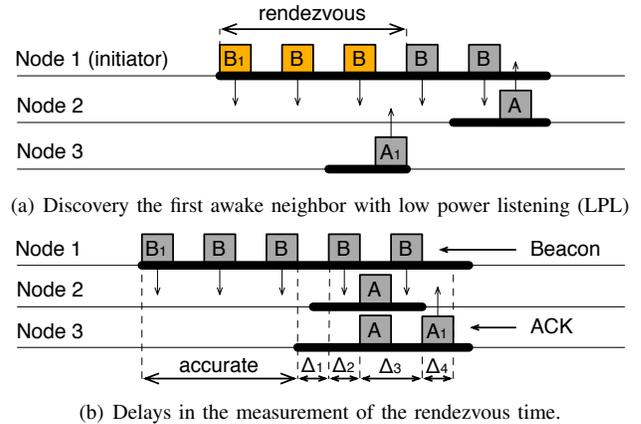


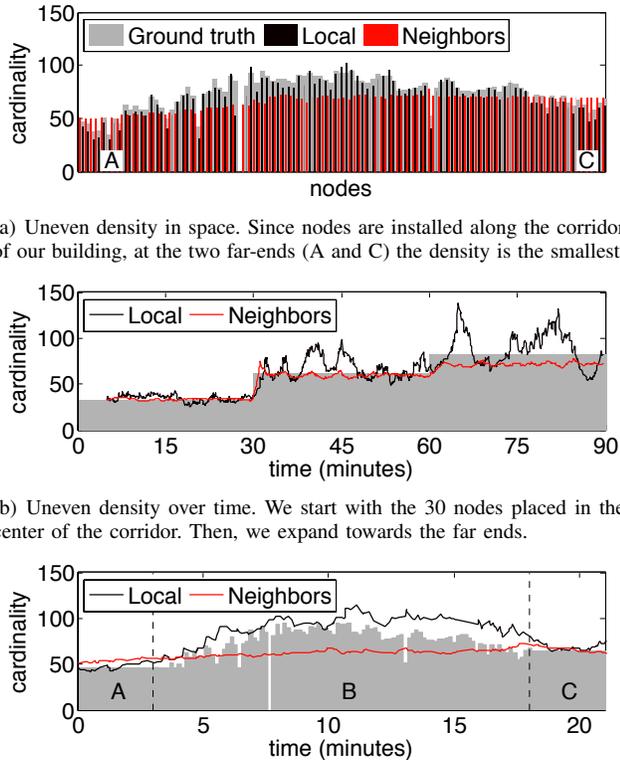
Fig. 1. Accurately measure the rendezvous time is essential for our estimator.

node, called *initiator*, wants to communicate (see Figure 1(a)), it starts sending a strobe of beacons (B). During this strobe, the initiator announces that the first node has not yet been discovered by setting a flag in its beacons (orange packets). When the *first* neighbor wakes up (node 3), it receives a flagged beacon and sends an acknowledgement (A1). After receiving A1, the initiator clears the flag and continues its strobe normally until the intended destination (node 2) responds.

Measure the rendezvous. In principle, the rendezvous time can be obtained by measuring the time elapsed between the first beacon (B1) and the first acknowledgement (A1). Unfortunately, as Figure 1(b) shows, the rendezvous phase includes different kinds of delays; namely, collisions (Δ_3), the transmission time of radio packets (Δ_2 , Δ_4), and other delays that are characteristic of the MAC protocol (Δ_1). To obtain an accurate measurement, in our mechanism the receiver (node 3) piggybacks the time elapsed since it woke up ($\Delta_1 + \Delta_2 + \Delta_3$) on its acknowledgment (A1).

Efficient collection. As any other estimator based on order statistics, the larger the number of samples, the closer the mean gets to the expected value and the more accurate the estimation. Unfortunately, in mobile networks nodes only have a limited amount of time to capture as many samples as possible. Thus, the number of samples w used to compute the average trades off accuracy for adaptability: the bigger w , the more time it takes to adapt to changes in cardinality. For this reason, whenever two nodes rendezvous, they exchange the average of their local samples. This exploits the spatial correlation on nodes' density and drastically increases the

*Marco Cattani was supported by the Dutch national program COMMIT/



(a) Uneven density in space. Since nodes are installed along the corridor of our building, at the two far-ends (A and C) the density is the smallest.

(b) Uneven density over time. We start with the 30 nodes placed in the center of the corridor. Then, we expand towards the far ends.

(c) Three mobile nodes move from one end of the corridor to the other.

Fig. 2. Cardinality estimation in different dynamic scenarios.

number of samples collected in time. Note that this is only possible because our estimator is designed to allow *all* nodes to perform concurrent estimations.

IV. EVALUATION

To evaluate our estimator, we ran a set of experiments on a testbed consisting of 100 nodes equipped with an MSP430 processor and a CC1101 radio [3]. For each experiment, we compute the ground-truth cardinality of a node as the number of neighbors that successfully exchanged a message at least once during the duration of an experimental run. With regards to the parameters, unless stated otherwise, our experiments use the following values: wakeup period (t_w) = sampling period = 1 s, window size (w) = 50, network size (n) = 100 nodes. Note that in the experiments, we differentiate between estimations computed using only local samples (*local*) and using neighbors' samples (*neighbors*) to highlight the different characteristics of the two.

Uneven density in space. For most cardinalities, our estimator has a relative error of $\approx 10\%$ for *local* and under 5% for *neighbors*. The main exception occurs at $n = 100$, when uneven densities among nodes in the testbed badly affect the spatial averaging mechanism, see Figure 2(a). For *local*, each node has a precise view of its own neighborhood, as shown by the accurate mapping between the individual node estimations and the ground truth. For *neighbors* however, by averaging the views of its neighbors, a node with a cardinality lower than its neighbors will overestimate its density, and vice versa.

Uneven density over time. Figure 2(b) shows a series of 90-minute experiments in which the neighborhood cardinality grows by 30 nodes each 30 minutes. With the same w , *neighbors* has a lower variance (better accuracy) and takes less time to adapt to changes (≈ 1 minute). During this time, a walking person (1 m/s) can cover approximately 60 meters while sampling the current neighborhood. Assuming a device with a transmission range of 50 m, our estimator should be able to cope with the dynamics of our application scenarios (crowd monitoring in open-air festivals).

Mobile nodes. As a final experiment, we equipped 3 colleagues with a sensor node and asked them to move according to a predefined path, see Figure 2(c). The experiment lasted 20 minutes. In the first 3 minutes, we asked our colleagues to stand on one far-end of the testbed. Subsequently, they slowly moved to the other end of the testbed (C). The slow movement (section B) was required to get an accurate measurement of the ground truth: at each step we waited approximately 10 s to obtain a reasonably accurate snapshot of the cardinality of the testbed's node that was closest to the mobile node. Figure 2(c) shows the estimated neighborhood cardinality of one of the mobile nodes and highlights the tradeoff between *local* and *neighbors*. If a quick and rough estimation is required, *neighbors* is the best solution. On the other hand, if a more accurate, but longer, measurement is needed *local* should be used.

V. CONCLUSIONS

In this poster we address the issue of determining the neighborhood cardinality in extreme wireless sensor networks, where node mobility necessitate a robust and agile approach. Moreover, we support a wide range of densities and concurrent estimations, requirements that are necessary in public safety applications such as crowd monitoring. Since traditional approaches cannot meet these stringent requirements, we developed an estimator based on observing the inter-arrival times of nodes waking up. We implemented our estimator on the Con-tiki OS, and evaluated it on a testbed with node densities up to 100 nodes. Our estimator achieves solid performance results with typical estimation errors below 10%, which compares favorably to state-of-the-art solutions. The estimator was also demonstrated to handle abrupt changes in density, exemplified through a few nodes moving through our testbed.

REFERENCES

- [1] P. Kannan and S. Venkatagiri. Low cost crowd counting using audio tones. In *Embedded Network Sensor Systems*, pages 155–168, 2012.
- [2] J. Weppner and P. Lukowicz. Collaborative Crowd Density Estimation with Mobile Phones. In *Embedded Networked Sensor Systems*, 2011.
- [3] M. Woehrle, M. C. Bor, and K. G. Langendoen. 868 MHz: a noiseless environment, but no free lunch for protocol design. In *Networked Sensing Systems*, INSS, pages 1–8, 2012.
- [4] W. Zeng, A. Arora, and K. Srinivasan. Low power counting via collaborative wireless communications. In *Information processing in sensor networks (IPSN)*, pages 43–54, 2013.

Poster Abstract: Synchronizing Trickle Intervals

Joakim Eriksson
Swedish Institute of Computer Science
Kista, Stockholm, Sweden
Email: joakime@sics.se

Omprakash Gnawali
Computer Science Department
University of Houston
Email: gnawali@cs.uh.edu

Abstract—Routing protocols for low-power wireless networks, such as CTP and RPL, use explicit control traffic between nodes to find neighbors, construct and maintain routing paths. To conserve power and reduce network congestion, the control traffic frequency must be kept low. Reducing control traffic will not only reduce power consumption, but also affect the quality of the routing paths constructed by the routing protocols, in terms both of delivery ratio and power consumption. Several existing protocols use control traffic reduction mechanisms such as Trickle. Trickle uses two mechanisms to achieve efficient routing: (1) rate adaption where the nodes send control traffic less often when the network is stable, and (2) suppression where a node avoids sending control traffic if the information has already been recently sent by neighboring nodes. In this work, we present scenarios that cause these mechanisms to result in pathology in control packet timing. We also present a modification to Trickle to address the pathology.

I. INTRODUCTION

Routing protocols for low-power wireless networks send control traffic for topology maintenance. Examples of such control traffic are neighbor reachability information [5] and routing beacons with path cost [1], [2]. Control traffic must reach all neighbors and is therefore sent using broadcast. In duty cycled low-power wireless networks, however, broadcast is significantly more power consuming and requires significantly more bandwidth than unicast transmissions.

Because of the expensive broadcast transmissions, routing protocols for low-power wireless face a fundamental trade-off: by sending more control traffic, the quality of the routing graph may be increased, but excessive control traffic significantly reduces network lifetime and also results in network congestion.

Many routing protocols use mechanisms such as Trickle [3] to minimize the number of control packets in the network while maintaining certain agility to network dynamics. Routing protocols such as CTP [2] and RPL [4], [5] use Trickle to time their control packets: decrease the control traffic interval when the routing cost change little or the network is not dynamic. They also suppress transmissions of control traffic when other nodes have already transmitted the same message.

The problem has an inherent complexity caused by the intricate interactions between the behavior of the radio medium, duty cycling mechanisms, the control traffic rate, and the data traffic. For example, although a routing protocol may create better routing paths by rapid dissemination of large number of control packets, the control traffic will cause a significant

congestion that interferes with the data traffic, causing lower application performance.

In this work, we describe the pathologies caused by the use of Trickle as mechanism for keeping control traffic at a low rate at stable topologies while maintaining a fast adaptation to changes. We also offer possible solutions to address these pathologies.

II. THE TRICKLE DISSEMINATION PROTOCOL

Trickle was originally designed for data dissemination in a multi-hop network. The data could be software version, or similar slow-changing data which is an ideal case for Trickle with its increasing intervals and suppression of already seen data. Later, Trickle was also used for timing the routing beacons. The Trickle protocol is specified by the pseudo-code in figure 1.

τ expires	\rightarrow	$\tau = \min(\tau \times 2, \tau_{max}), c = 0, t = [\frac{\tau}{2}, \tau]$
t expires $\wedge c < k$	\rightarrow	transmit M_i
receive M_i	\rightarrow	$c = c + 1$
receive $M_{j>i}$	\rightarrow	$\tau = \tau_{start}, c = 0, t = [\frac{\tau}{2} - \tau]$

Fig. 1. The rules of the Trickle mechanism. When the maintenance interval (τ) expires, τ is doubled if not already at its maximum (τ_{max}). The redundancy counter c is also set to zero. When t expires, a metadata message is sent if c is less than k . When an incoming metadata message is received, c is increased if the message is consistent with the current metadata; otherwise the timers and intervals are reset.

Trickle disseminates new information fast, sends very little information when there is no new information available for dissemination, and does load balancing between the nodes.

III. TRICKLE TRICKED

When all nodes have synchronized intervals, i.e., Trickle timers are approximately aligned across the nodes, Trickle handles propagation of meta-data efficiently and load balances the propagation between the nodes in the network. Unfortunately, the load balancing fails when the nodes start to lose synchronization with each other. In a network with bottleneck topology and desynchronized Trickle timer, information propagation can be slow.

One case where Trickle load balancing fails is when there is a large network of synchronized nodes and a new node joins by listening to the network beacons. The first beacon will be sent very close to half the interval which will cause this node to be unsynchronized with the other nodes. Figure 3

shows the consequence of having an unsynchronized node in a single-hop network consisting of 10 nodes.

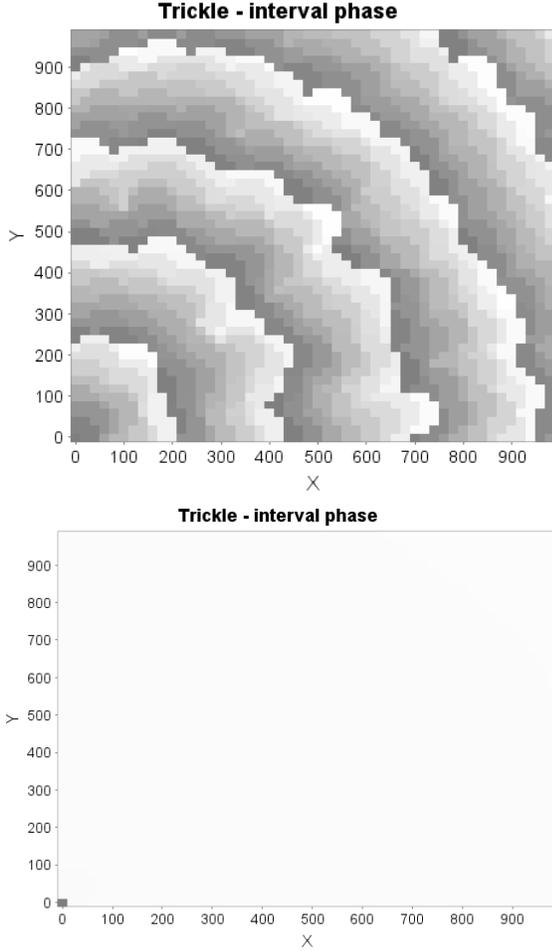


Fig. 2. Results from a simulation in a 2500-node network. On the top, we run the unmodified Trickle. We use the node in the lower left corner as reference to visualize the timer phase. White corresponds to nodes that have the same interval phase as the reference node. The darker the color, the more out of sync the interval. On the bottom figure, we run the modified Trickle that resynchronizes its interval.

IV. TRICKLE WITH RESYNCHRONIZATION OF INTERVALS

The desynchronization in Trickle can be eliminated by adding a synchronization mechanism that resynchronizes the time intervals in Trickle. This mechanism makes all nodes execute their Trickle listen and transmission intervals at the same time which improves the efficiency and also removes any lack of load balancing that may appear in a network. The algorithm described in Figure 4 achieves this goal.

The efficiency of the modified Trickle at a steady state when the intervals have been synchronized will approach “perfect” (e.g. k packet per interval in a network with full reachability) and in the worst case, assuming that the synchronization fails completely, it will be as efficient as Trickle with an interval of $\tau - w_s \frac{1}{4} \tau$. By resynchronizing, the Trickle protocol becomes both more efficient and better load balanced which reduces

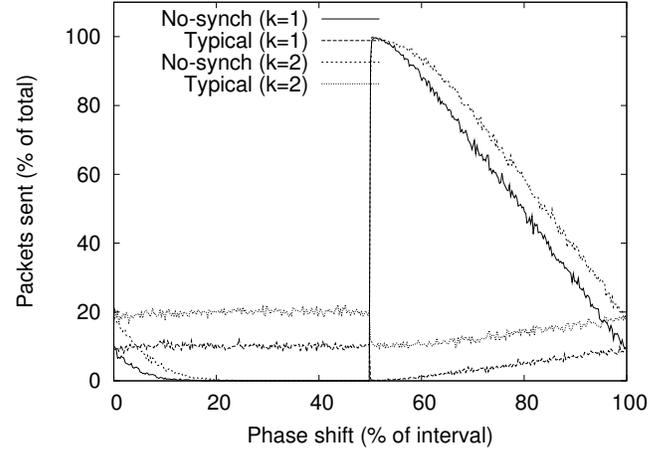


Fig. 3. Results from a simulation of a 11-node network. The network is setup with 10 synchronized nodes and one node that is out of phase with the other 10. k is set to 2. The phase is shifted from 0% to 100% of the maximum interval τ_{max} . The load balancing of Trickle breaks quite dramatically around 50% phase shift. Some nodes end up sending packets almost every interval while one did not perform any transmission.

τ_s expires	\rightarrow	$\tau = \min(2\tau, \tau_{max}), \tau_s = \tau - \delta_s,$
		$c, \delta_s = 0, t = \lceil \frac{\tau_s}{2}, \tau_s \rceil$
t expires, $c < k$	\rightarrow	transmit M_i
receive M_i	\rightarrow	$c = c + 1$
receive $M_i, T \in [\frac{\tau}{4}, \frac{\tau}{2}]$	\rightarrow	$\delta_s = w_s(T - \frac{\tau}{2})$
receive $M_{j>i}$	\rightarrow	$\tau = \tau_{start}, c = 0, t = \lceil \tau/2 - \tau \rceil$

Fig. 4. Trickle modified for resynchronization of intervals. The idea is to start the next interval earlier if this node got a Trickle message during its second half of its listen only period.

the propagation delays in bottleneck topologies. Figure 2 compares the synchronization properties of the original and the modified Trickle. We find that the modified Trickle was able to synchronize the schedules of the nodes across the network while the original Trickle was not.

V. ACKNOWLEDGEMENTS

We would like to thank Philip Levis and others in the IETF ROLL working group for the active discussions related to Trickle and its use in RPL.

REFERENCES

- [1] D. De Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *Proceedings of the International Conference on Mobile Computing and Networking (ACM MobiCom)*, pages 134–146, San Diego, CA, USA, 2003. ACM.
- [2] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis. Collection tree protocol. In *Proceedings of the International Conference on Embedded Networked Sensor Systems (ACM SenSys)*, Berkeley, CA, USA, 2009.
- [3] P. Levis, N. Patel, D. Culler, and S. Shenker. Trickle: A self-regulating algorithm for code propagation and maintenance in wireless sensor networks. In *Proceedings of the USENIX Symposium on Networked Systems Design & Implementation (NSDI)*, March 2004.
- [4] J.P. Vasseur and A. Dunkels. *Interconnecting Smart Objects with IP: The Next Internet*. Morgan Kaufmann, 2010.
- [5] T. Winter (Ed.), P. Thubert (Ed.), and RPL Author Team. RPL: IPv6 Routing Protocol for Low power and Lossy Networks. RFC 6550.

Poster Abstract: The Purpose and Benefit of Segmenting a Multi-Transmitter Network

Douglas Carlson
Computer Science Department
Johns Hopkins University
carlson@cs.jhu.edu

Omprakash Gnawali
Computer Science Department
University of Houston
gnawali@cs.uh.edu

Abstract—In a multi-transmitter network, multiple wireless nodes transmit data at the same time. If these transmissions are timed close enough to each other, they interfere non-destructively and can be correctly decoded. Systems such as CX, Glossy, and LWB use this principle to achieve efficient and reliable network-wide synchronization or communication. In practice, many wireless sensor network deployments are “patchy”, i.e., a few widely-spread areas are densely instrumented. While prior work in multi-transmitter networks relied on a single, globally-coordinated network, intuition tells us that subdividing the network will make it easier and more energy-efficient to coordinate schedules. In this work, we show how multi-transmitter networking can be applied to segmented networks and evaluate its effect on energy consumption and packet delivery rates.

I. INTRODUCTION

Glossy [1], Low-power Wireless Bus (LWB) [2], and CX [3] are examples of an emerging family of multihop wireless sensor network (WSN) communication protocols which we refer to as *multi-transmitter networking*. These systems leverage a combination of non-destructive concurrent transmissions and radio capture effect to perform fast network floods that reach all nodes with high probability: every node receiving a packet rebroadcasts it at precisely the same time, reducing destructive interference. In addition to the high yield, good throughput, and low energy consumption demonstrated in LWB, these methods require little routing state to work. This makes them suitable for networks with high degrees of node mobility, and may help in challenging environments where existing routing methods struggle to find high quality links. These systems rely on a single globally-coordinated TDMA schedule to prevent conflict and to enable duty-cycling.

In practice, many WSN deployments for environmental monitoring consist of a set of densely-instrumented patches. Each such patch comprises multiple *leaf* nodes (which measure their environment) and one *router* node (which enables inter-patch communication). Many seminal sensor network deployments have employed this network architecture in some form. In this work, we wish to adapt the globally-coordinated multi-transmitter networking approach to one that is well-suited to this common deployment pattern.

In patchy networks, most nodes are leaves. Adding a leaf to the network adds both coordination overhead (in terms of discovering that node and assigning it time in the schedule) and forwarding load to other devices in the network. Our

goal is to restrict this impact to just the patch where the new node was added. Each patch is headed by a router, which is responsible for autonomously retrieving the data from its patch. These routers work together to transfer their collected data to a central collection point. In this work, we experimentally verify and quantify the reduced energy cost at leaf nodes and added energy cost at router nodes due to a patchy organization compared to a flat network.

II. CX: A MULTI-TRANSMITTER NETWORK

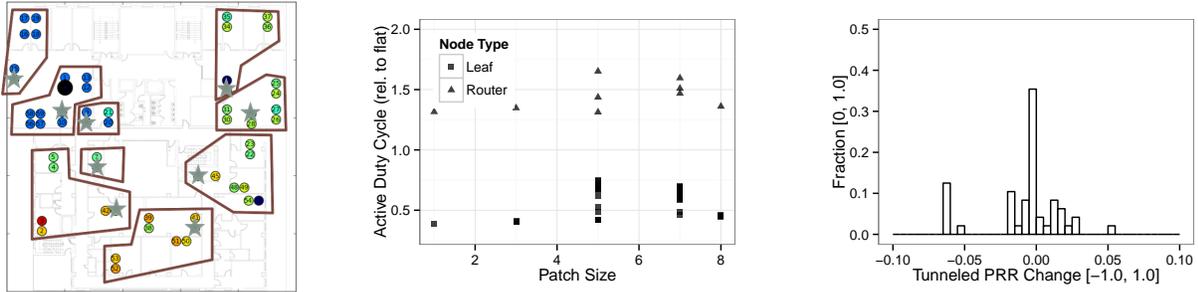
CX is an example of a multi-transmitter network system [3]. In CX, all communication takes place in the form of multi-transmitter floods. Through the use of a hop counter in each packet, nodes learn their relative distances to each other. The hop count information allows nodes to estimate whether they are between a source and destination for a given transfer. These measurements are the basis for our forwarder-selection method, which keeps nodes between the source and destination active while allowing the rest of the network to sleep.

Results from our 66-node indoor testbed show that forwarder-selection reduces duty cycle by 30% on average over simple concurrent flooding while maintaining an average packet reception ratio (PRR) of 99.4%. In the same setting, the average node throughput increases by 49% over simple flooding by using the collected distance information to tighten inter-packet spacing.

III. SEGMENTING MULTI-TRANSMITTER NETWORKS

We subdivide the full network into multiple patches at deployment time. The nodes in different patches are assigned different radio channels through the use of a Graphical User Interface (GUI) tool based on the user’s knowledge of how the nodes are physically deployed. We designate some nodes as leaves and designate others as routers, where routers independently collect data from the leaves in their patch. A basestation node periodically downloads the collected data from each router while the leaves keep their radios off to save energy. We use the technique of CXFS [3] to perform each of these data collection steps reliably and efficiently.

Figure 1a shows an example of such a network. Each download, whether from leaves to a router or from routers to a basestation, follows the same basic pattern.



(a) Segmented 50 m x 50 m indoor testbed. Stars mark router nodes, and the black circle marks the ultimate data sink. (b) Duty cycle increases at leaf with larger patches. 1.0 on the Y-axis indicates no change, higher values indicate worse duty cycle under segmentation than in a flat network. (c) End-to-end PRR changes for nodes under segmentation.

Fig. 1: Testbed layout and experiment results.

All nodes use Low-Power Probing [4] to coordinate their wake-up/sleep cycles. Once a network segment (e.g. basestation and all routers or a single router and the leaves in its patch) is active, the download proceeds in a series of *slots*, each of which grants exclusive access to a single source node for a period of time (on the order of seconds). The sink sends a Slot Assignment message to one of its immediate neighbors, which responds with a Status message. This Status message informs the sink of the assignee’s one-hop neighborhood and carries the necessary information to perform forwarder selection. This allows the sink to discover the members of its network segment. During its slot, a node sends any outstanding data it may have and ends with a message indicating whether or not it still has data pending. When the sink determines that no more nodes have outstanding data, it stops assigning slots and the network segment returns to the idle state.

By assigning each patch to a separate radio channel, we allow these collections to take place in parallel and without interference. Routers perform a download from their patch immediately following a download by the basestation.

IV. EXPERIMENT SETUP

We implemented segmented CX on a CC430-based platform [5] on our indoor testbed. Prior to segmenting the network, the mean end-to-end PRR is above 99.5% from leaf to root (with and without forwarder selection), while the root to leaf PRR is 98.7%. This low PRR appears to be due to poor connectivity at a single node. We divided our testbed into 9 distinct patches with 1-8 leaf nodes in each patch.

The sink collects 75 packets per download, with 100-byte payloads, which is roughly equivalent to the daily data rate for our target application (8 sensors per node, 2-byte samples, 10-minute sampling rate, and associated timing information and metadata). Transmission power was set to -6 dBm at all nodes. Our platform can be assembled with a radio amplifier which may be used to extend the distance between patches.

V. RESULTS

Figure 1b plots the duty cycle changes in the network. The average leaf duty cycle drop is 56% of its single-tier level when they are grouped into patches. The improvement in leaf

node duty cycle is not free. While routers enjoy the same short downloads that the leaves in their patch do, they also have to retransmit all of this data to the root and suffer an additional wakeup. Their duty cycle is 50% higher than it is in a single-tier network, on average. That being said, the total energy consumption of the network decreases under segmentation, consuming only 70.8% as much as the flat network.

The end-to-end packet reception ratio for leaf nodes may change when the network is segmented. Figure 1c shows the distribution of the changes in PRR experienced by the leaf nodes when moving from a flat network to a tiered network. The average packet reception ratio drops to 96.91%, primarily due to the poor router-to-root PRR of a single router. 22 out of the 57 leaf nodes see improvements ranging up to 3.5%. This can occur if their router has a good PRR to the root, and their in-patch PRR is better than their flat-network PRR.

VI. CONCLUSIONS

In this work, we adapt multi-transmitter networks to take advantage of the patchy layout present in many wireless sensor network deployments. We found that segmenting a multi-transmitter networks into patches results in significant energy savings for the leaves at the expense of higher energy expenditure on the routers. Sending packets through the routers in a patchy networks is limited by the achievable PRR to the quality of connectivity between the routers and the downstream nodes.

Acknowledgment: Partial support through NSF-IIS-1111507.

REFERENCES

- [1] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh, “Efficient network flooding and time synchronization with glossy,” in *IPSN*, 2011, pp. 73–84.
- [2] F. Ferrari, M. Zimmerling, L. Mottola, and L. Thiele, “Low-power wireless bus,” in *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*, ser. SenSys ’12. ACM, 2012, pp. 1–14.
- [3] D. Carlson, M. Chang, A. Terzis, Y. Chen, and O. Gnawali, “Forwarder Selection in Multi-Transmitter Networks,” in *Proceedings of the 9th International Conference on Distributed Computing in Sensor Systems (DCOSS 2013)*, May 2013.
- [4] R. Musłoiu-E., C.-J. Liang, and A. Terzis, “Koala: Ultra-low power data retrieval in wireless sensor networks,” in *Proceedings of the 7th international symposium on information processing in sensor networks (IPSN)*, April 2008, pp. 421–432.
- [5] Texas Instruments, “MSP430 SoC with RF Core,” Available at <http://www.ti.com/product/cc430f5137/>, 2010.