

Exploring Sensor Network Communication Strategies with the Mote-in-the-Loop Approach

Minyan Hong, Erik Björnemo
Uppsala University
Uppsala, Sweden

minyan.hong@gmail.com, Erik.Bjornemo@angstrom.uu.se

Thiemo Voigt
Swedish Institute of Computer Science (SICS)
Kista, Sweden
thiemo@sics.se

Abstract—The performance of wireless sensor networks is heavily influenced by the surrounding environment. Hence, in different scenarios different communication strategies, i.e., different packet sizes or retransmission strategies etc., are preferable. There are, however, currently no appropriate methods to investigate the optimal strategy. Analysis and simulation rely on abstract, often unrealistic assumptions and experiments with real hardware are time-consuming. Towards this end, we present *mote-in-the-loop*, a new approach to explore sensor network communication strategies. Our experimental results demonstrate the feasibility of our approach and make us believe that *mote-in-the-loop* can become a powerful and useful tool.

Index Terms—sensor networks, performance, signal generation, interference

I. INTRODUCTION

Sensor networks are being deployed in a range of different environments, such as industry plants, rainforests and offices. Each environment has its own characteristics [7], [17] which might be impacted by the environmental conditions [1], [3], [6]. Hence, the most suitable communication strategy will differ accordingly – packet sizes, retransmission schemes, error correcting codes, etc.

It is, however, difficult to investigate the most appropriate communication strategies for the environment of an intended deployment. On the one hand, simulations are seldom realistic enough as they do not model the environment in every intricate detail. On the other hand, real-world experiments with deployed nodes are important but time-consuming, difficult to repeat, and to some extent dependent on hardware and software [13]. For example, a bug in the software might make measurements collected during an extensive time useless [15]. We need an easier way of testing which still captures realistic communication environments and provides repeatability.

We propose a new approach to investigate communication strategies. Our approach uses a combination of on-site radio channel and interference measurements, real sensor network hardware as well as a signal analyser and a signal generator. The advantage of our approach is that once the channel measurements are made, we have a deterministic and repeatable way of investigating the most suitable communication strategy in the lab and for different sensor node hardware. Additionally, we can quickly test new hardware and new implementations by simply recording new packets.

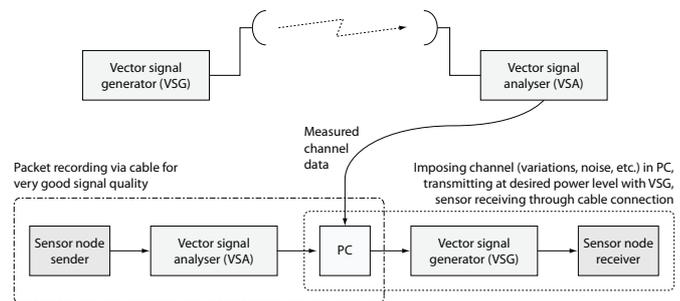


Fig. 1. Experimental setup for repeatable testing of communication strategies using real channel data.

We present our basic approach with results that are partly peculiar but demonstrate the correctness of our approach since they demonstrate hardware properties that other researchers also have discovered. Then we compare the results of the *mote-in-the-loop* approach with results from real measurements with microwave oven interference. The experiments demonstrate the feasibility of our approach.

The rest of the paper continues as follows. First, we discuss our approach in more detail in Section II. In the following two sections we present experimental results that validate the correctness of our approach. After discussing related work in Section V we present our conclusions.

II. APPROACH

In this section we describe our approach for the investigation of communication strategies.

A. Basic Setup

The setup consists of two motes, a vector signal analyser (VSA) and a vector signal generator (VSG)¹, see Figure 1.

A modern vector signal analyser/generator is an advanced instrument with the following typical characteristics: Large frequency range; Large signal bandwidth; Accurate power reading/setting. These features render the instrument flexible and facilitate tests outside the reach of mote-to-mote communication such as the transmission of recorded signals at very precisely set power levels.

¹In our case the 2810 VSA and the 2910 VSG from Keithley.

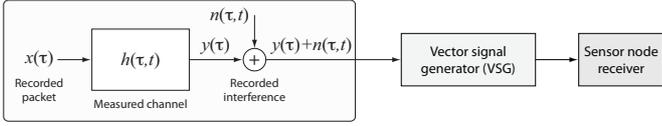


Fig. 2. Introduction of fading channel and interference. The packet $x(\tau)$, the channel $h(\tau, t)$ and the interference $n(\tau, t)$ are all complex-valued to contain both amplitude and phase information. The variable t shows that the channel and the interference can have time-varying characteristics.

The motes are TmoteSky sensor nodes [19] that feature a CC2420 radio and run the Contiki operating system [10]. The sending mote sends one or more packets that the signal analyser resolves to in-phase and quadrature (IQ) values that can be stored on the PC, for example, using Matlab. The radio communication between mote and signal analyser is via a cable to avoid external interference and achieve large signal-to-noise ratio (> 60 dB). With software on the PC, we can instruct the signal generator to replay the packets received by the signal analyser and transmit them to the mote as depicted in Figure 1. We can further vary the output power of the signal generator, e.g. according to measured channel gains. In particular, we are able to collect measured channel data from different environments to emulate the impact of the environment on communication. This way, we expect to be able to find communication strategies tailored to the environment.

At the receiving mote we can measure e.g. packet reception rate but also retrieve the received signal strength indicator (RSSI), the link quality indicator (LQI) and noise floor values from the on-board radio.

B. Including Fading and Interference

In Figure 2 we depict the general procedure for including fading and interference. Note that there is a choice when it comes to the thermal noise as it can either be introduced artificially in the PC, as part of $n(\tau, t)$, or by using the real receiver noise and a scaled output power from the VSG.

III. EVALUATION AND PROOF OF CONCEPT

In this section we provide some experiments that validate that our approach produces meaningful results.

A. Basic RSSI experiment

We verify the CC2420's RSSI readings by repeatedly replaying a recorded packet at increasing power levels. Figure 3 shows the results over the power range in which the mote actually receives the packets and can hence measure and report the RSSI (down to approximately -95 dBm). The figure shows the expected overall linear relationship, with a small variance in the RSSI readings. However, we specifically note two regions – at output powers of -40 dBm and -25 dBm – where the linear relationship between RSSI and VSG output power is disturbed and the sample variance is larger. This reflects an inaccuracy of in the RSSI reading mechanism that also Chen and Terzis have observed [8]. Note that this inaccuracy thus *confirms* the correctness of our approach.

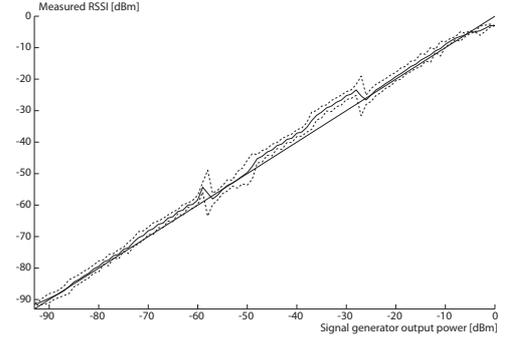


Fig. 3. RSSI from the CC2420 as a function of the VSG signal output power. The mean ± 3 standard deviations are given together with the ideal straight-line response. The curves are based on 10000 RSSI readings per power setting.

B. Basic PER in a Gaussian channel

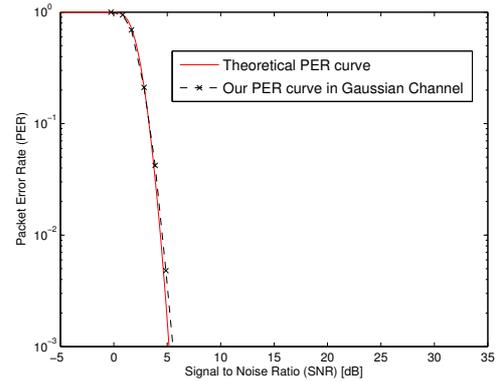


Fig. 4. Packet error rate (PER) in a Gaussian channel. Theoretical and measured results match very well. changes from our test in no fading channel, with the theoretical PER curve as a reference.

Figure 4 illustrates how packet error rate (PER) and signal to noise ratio (SNR) are related in our Gaussian channel. Here the SNR is the average signal to noise ratio based on collected average RSSI and average noise. The theoretical packet error rate is calculated from the IEEE 802.15.4 packet specification using a 4-byte payload and a one byte address. The minimum Hamming distance of 12 between the chip sequences is used; we assume that six chip errors per spreading code word are correctable while seven or more lead to a packet error. This results in an upper bound on the packet error rate [14].

Figure 4 shows that the packet error rate (PER) decreases as the SNR increases, and there is a sudden drop at a SNR level less than 5 dB. All data following have a PER of zero. The solid curve in Figure 4 is the theoretical PER curve corresponding to our packet length [14]. The graph demonstrates that it matches our results very well which shows that effectiveness of our approach.

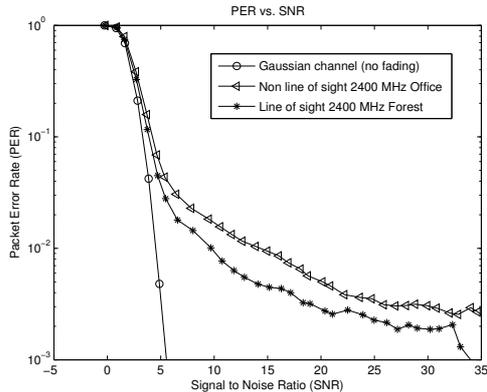


Fig. 5. Packet error rate (PER) for a range of signal-to-noise ratios (SNR). The channel was block fading, that is to say roughly constant during a packet but changing on an inter-packet time scale.

C. Repeatable communication in fading channels

While real world deployments in one respect constitute the ultimate test of a sensor network and its communication strategies, it can be very difficult to compare results for different deployments at different times. One reason is that variations in link quality – channel fading – are different at different times and locations. During research and development, it is therefore desirable to have a repeatable approach which still is much more realistic than simulation. Our proposed approach is a step in this direction, and we here show an example of the results we can achieve using real-world channel data. The channel characteristics used here consists of traces that we have collected in office and forest environments [2].

We used the approach depicted in Figure 2 to study the channel impact without interference for illustrative purposes. By the use of 10000 packets for each received average signal-to-noise ratio, we obtain packet error rate curves for three cases: No fading, measured office fading and measured forest fading. The channel data is applied so that block fading is achieved, that is a fairly constant channel during packet transmissions.

The results in Figure 5 show how the fading introduces error floors, starting at packet error rates of around 3 percent. The difference between the fading channels is not as extreme as one might expect, but it should be noted that the terms "line of sight" and "non line of sight" are inadequate to describe the difference. In fact, the office setting allowed some penetration through walls which resulted in "partial line of sight" (non-Rayleigh fading). Additionally, the forest setting was not pure line of sight because of the antennas being very close to the ground.

IV. EXPERIMENTAL EVALUATION: MICROWAVE INTERFERENCE

To evaluate our approach we perform experiments where we compare our mote-in-the-loop approach with microwave interference on real motes. The goal of the experiment is to

investigate the realism of the results achieved with the mote-in-the-loop approach.

A. Background: Microwave Interference

Microwave ovens are a source of interference as they, similar to most low power radios, also operate in the 2.4 GHz band. We have used Wi-Spy to collect the noise emitted by microwave ovens at a distance of 1 meter.

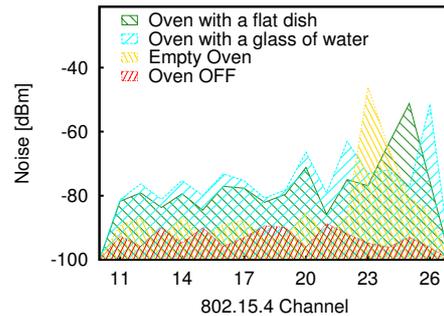


Fig. 6. Spectral characteristics of microwave interference

Figure 6 depicts the spectral characteristics of the interference, i.e., the noise in different channels. Measurements with different microwave oven models and different content show that the interference is strongest on the channels between 20 and 26 but also depend on what is in the oven.

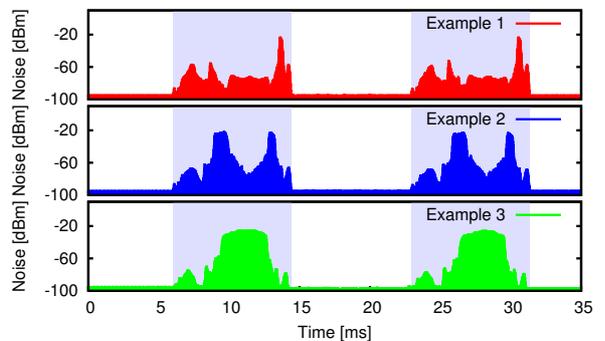


Fig. 7. Temporal characteristics of microwave interference

Figure 7 depicts the temporal characteristics of microwave interference. The interference follows an on-off pattern but also inhibits a large degree of burstiness. The figure suggests that when a radio device is close to a microwave, at roughly half of the time packet reception could be impossible due to the interference. When moving further away from the microwave the interference will eventually decrease to zero.

B. Experimental approach and setup

In order to investigate the realism of our approach, we perform a set of experiments. First, we set up an experiment where a microwave is turned on while two motes close

to the microwave are communicating. Second, we sample the interference $n(\tau, t)$ (see Figure 2) from the microwave. Third, we perform the mote-in-the-loop approach that uses the sampled interference $n(\tau, t)$. We perform the first and second experiment in an anechoic chamber.

In the first experiment, we use two Tmote Sky nodes, one sender and one receiver. In this experiment, we change the transmit power of the sender node to get different signal to interference ratio (SIR). We sent 10000 packets for each transmit power level within 60 seconds and ensure that the microwave is always turned on while the sensor nodes are communicating. As in the other experiments we measure the PER.

In the second experiment we record a trace using the signal analyzer. We fix the position of the antenna and use a cable to connect the antenna and the signal analyzer that we place outside the anechoic chamber to avoid any impact on the signal propagation. The signal analyzer will record five seconds of interference each run, then resolve the measurements to I/Q values and load these data into PC using MATLAB. We record only five seconds of interference each run because of the memory limitation of the signal analyzer. As above, we put the cup of water into the microwave oven when collecting data. We record traces for several frequencies and several trigger levels that make the signal analyzer start recording.

For our mote-in-the-loop experiments, we use the following setup. The signal generator generates the recorded interference from the experiment above. Two sensor nodes are programed to be sender node and receiver node respectively. As in the first experiment the sender node sends 200 packets per second, i.e., it will transmit 1000 packets in 5 s, i.e., one interference sequence covers 1000 packets. We use a splitter/combiner to combine packets from the sender node and the interfering signal from the signal generator, i.e., we perform superimposing. The combiner sends the superimposed signal to the receiver node. All connections from/to the combiner are wired in order to reduce environment influence as much as possible.

C. Results

Figure 8 and Figure 9 show the results of the experiment with real sensor nodes and our mote-in-the-loop approach for two different IEEE 802.15.4 channels, namely channel 24 at 2470 MHz and channel 25 at 2475 MHz. For the mote-in-the-loop approach, we compute the average over all runs that used different noise traces for different trigger values.

The figures show that the mote-in-the-loop approach matches the results of the real experiments quite well. There are, however, some differences that we explain below. Note also that although the tendencies in Figure 8 and Figure 9 are similar there are slight differences. In the latter figure the packet error rate is higher. The reason for this is that microwave oven interference is different in different channels as shown in Figure 6.

In the two graphs, the curves show similar tendencies. They can be roughly divided into three parts: a steep part to the right and the left and a slowly decreasing part in the middle.

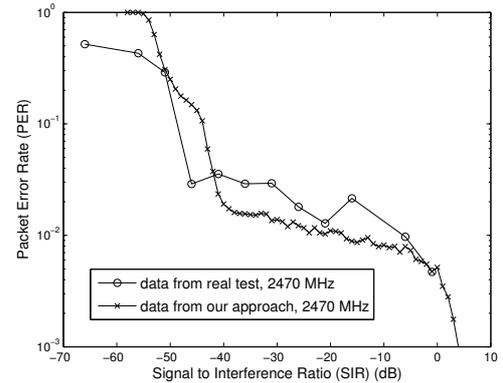


Fig. 8. PER under microwave oven Interference with frequency 2470 MHz

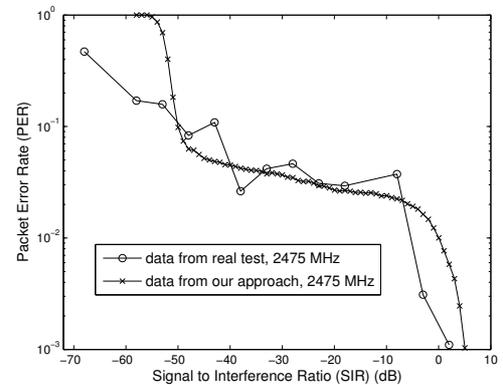


Fig. 9. PER under microwave oven Interference with frequency 2475 MHz

To the very right, we see the situation when there is no interference and hence the PER is zero. When we increase the SIR, the interference starts to corrupt packets and the packet error rate increases. In the real experiment the SIR is increased by increasing the output power level of the sensor node that is transmitting while in the mote-in-the-loop approach the SIR is increased by decreasing the power level of the signal generator that replays the recorded microwave noise.

In the middle part the SIR gradually increases to the left which causes a higher packet error rate. Note that the increase in SIR is similar to moving a receiving real sensor node closer to the microwave. Figure 7 has shown the burstiness of the microwave interference which explains the gradual increase in the number of corrupted packets.

In the very left part of the graphs we see that PER for the mote-in-the-loop approach approaches one for a SIR at about -50 to -60 dB whereas the experiment with real nodes approaches a PER of about 0.5. The reason for the latter is again explained in Figure 7 that demonstrates that a microwave emits signal at about 50% of the time which leads to a PER of about 0.5. In the mote-in-the-loop approach the thermal noise impacts the results and causes additional packet loss.

D. Limitations

As shown in the results above, there are two major limitations. One is the relatively low memory of the signal analyzer which does not allow us to capture long traces. Furthermore, the internal noise level of the signal analyzer contributes to uncertainties. While signal analyzers cover a large bandwidth, we have noticed that low power radios such as the CC2420 are actually more accurate at the smaller bandwidth they cover. Nevertheless, valuable insights can already be gained using our approach and we believe that both issues raised above are less problematic in newer versions of the equipment in particular the vector signal analyzer.

V. RELATED WORK

There exists a large number of simulation tools for sensor networks [12] that are developed for different purposes such as to ease development [16], [18] or to provide realistic wireless channel and radio models for algorithm development [5].

JamLab is a tool that augments existing sensor network testbeds with capabilities to generate realistic interference [4]. The authors present and evaluate implementations of models of different devices that are sources of interference in the 2.4 GHz band. Furthermore they provide a feature to record and replay interference. Our work is more generic in that we can manipulate the interference patterns in many different ways. While JamLab is a mote-only solution, our approach requires additional hardware in terms of signal generators.

Related to our work are also deployment tools, such as the one proposed by Ritter et al. [20] that help the network operators to deploy a network and a radio mapping tool for indoor environments [9]. Another deployment support tool is the Deployment Support Network (DSN) that provides a second wireless backbone network to observe and control the primary network after deployment [11]. Our approach is different in that we enable the optimization of various parameters before deployment.

VI. CONCLUSIONS AND FUTURE WORK

We have presented mote-in-the-loop, a new approach for communication strategy exploration. The approach allows to experiment with communication strategies that makes less assumptions on the real world than simulations and hence might be more realistic. Furthermore, the experiments are less tedious than experiments with real sensor nodes. Moreover, the experiments with the mote-in-the-loop approach are deterministic once the data has been collected. They also allow to analyze interesting behaviour in more detail. With some further extensions such as feedback from the sensor node to the VSG in order to trigger retransmission we believe that mote-in-the-loop will become a powerful and useful tool.

ACKNOWLEDGEMENTS

Thanks to Carlo Alberto Boano for the figures showing the microwave characteristics. This work was carried out within WISENET (Uppsala VINN Excellence Center for Wireless Sensor Networks) funded by VINNOVA, Uppsala University,

Banverket, CRL Sweden, FOI, Imego, JonDeTech, Pricer, SenseAir, SICS, TNT-Elektronik, TermoSense, VTT, Upwis, AAC Microtec, and WiseNet Holding. The work has been partially supported by CONET, the Cooperating Objects Network of Excellence, funded by the European Commission under FP7 with contract number FP7-2007-2-224053.

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 *Proceedings of the fifth Workshop on Embedded Networked Sensors (HotEmNets 08)*, June 2008.
- [2] Erik Björnmemo. *Energy Constrained Wireless Sensor Networks : Communication Principles and Sensing Aspects*. Institutionen för teknikvetenskaper, Uppsala University, 2009.
- [3] C.A. Boano, N. Tsiftes, T. Voigt, J. Brown, and U. Roedig. The Impact of Temperature on Outdoor Industrial Sensornet Applications. *Industrial Informatics, IEEE Transactions on*, (99):1–9, 2010.
- [4] C.A. Boano, T. Voigt, C. Noda, K. Römer, and M. Zúñiga. JamLab: Augmenting Sensornet Testbeds with Realistic and Controlled Interference Generation. In *ACM/IEEE IPSN'11*, Chicago, USA, April 2011.
- [5] A. Boulis. Castalia: revealing pitfalls in designing distributed algorithms in WSN. In *Proceedings of the 5th international conference on Embedded networked sensor systems*, Cork, Ireland, June 2007.
- [6] C. A. Boano and J. Brown and Z. He and U. Roedig and T. Voigt. Low-Power Radio Communication in Industrial Outdoor Deployments: The Impact of Weather Conditions and ATEX-compliance. In *Sensappeal*, Athens, Greece, September 2009.
- [7] M. Ceriotti, M. Chini, A. Murphy, G. P. Picco, F. Cagnacci, and B. Tolhurst. Motes in the Jungle: Lessons Learned from a Short-Term WSN Deployment in the Ecuador Cloud Forest. In *International Workshop on Real-World Wireless Sensor Networks*, Colombo, Sri Lanka, December 2010.
- [8] Y. Chen and A. Terzis. On the Mechanisms and Effects of Calibrating RSSI Measurements for 802.15. 4 Radios. In *European Conference on Wireless Sensor Networks*, Coimbra, Portugal, February 2010.
- [9] O. Chipara, G. Hackmann, C. Lu, W.D. Smart, and G.C. Roman. Practical modeling and prediction of radio coverage of indoor sensor networks. In *ACM/IEEE IPSN*, Stockholm, Sweden, 2010.
- [10] A. Dunkels, B. Grönvall, and T. Voigt. Contiki - a lightweight and flexible operating system for tiny networked sensors. In *Proceedings of EmNets-I*, Tampa, Florida, USA, November 2004.
- [11] M. Dyer, J. Beutel, T. Kalt, P. Oehen, L. Thiele, K. Martin, and P. Blum. Deployment support network a toolkit for the development of WSNs. In *Conference on Wireless sensor networks*, February 2007.
- [12] J. Eriksson. Detailed Simulation of Heterogeneous Wireless Sensor Networks. Licentiate thesis, May 2009.
- [13] Niclas Finne, Joakim Eriksson, Adam Dunkels, and Thiemo Voigt. Experiences from two sensor network deployments self-monitoring and self-configuration keys to success. In *WWIC 2008*, May 2008.
- [14] M. Hong. A Mote-in-the-Loop Approach for Exploring Communication Strategies for Sensor Networks. Master's thesis, Uppsala University, 2010.
- [15] K. Langendoen, A. Baggio, and O. Visser. Murphy loves potatoes: experiences from a pilot sensor network deployment in precision agriculture. In *Proceedings of IEEE IPDPS'06*, Rhodes Island, Greece, April 2006.
- [16] P. Levis, N. Lee, M. Welsh, and D. Culler. Tossim: accurate and scalable simulation of entire tinyos applications. In *Proceedings of the first international conference on Embedded networked sensor systems*, 2003.
- [17] L. Mottola, G.P. Picco, M. Ceriotti, S. Gunva, and A. Murphy. Not all wireless sensor networks are created equal: A comparative study on tunnels. *ACM Transactions on Sensor Networks*, 7(2):1–33, 2010.
- [18] F. Österlind, A. Dunkels, J. Eriksson, N. Finne, and T. Voigt. Cross-level sensor network simulation with cooja. In *Proceedings of the First IEEE International Workshop on Practical Issues in Building Sensor Network Applications (SenseApp 2006)*, Tampa, Florida, USA, November 2006.
- [19] J. Polastre, R. Szewczyk, and D. Culler. Telos: Enabling ultra-low power wireless research. In *Proc. IPSN/SPOTS'05*, Los Angeles, CA, USA, April 2005.
- [20] H. Ritter and S. Homfeld. Real-world wsn deployment tool. In *Demo Session of the ACM Workshop on Real-World Wireless Sensor Networks (REALWSN'06)*, Uppsala, Sweden, June 2006.