KonTest: A Wireless Sensor Network Testbed at Vrije Universiteit Amsterdam

Konrad Iwanicki,*† Albana Gaba,* and Maarten van Steen*
*Vrije Universiteit, Amsterdam, The Netherlands
†Development Laboratories (DevLab), Eindhoven, The Netherlands
{iwanicki, agaba, steen}@few.vu.nl

Technical Report IR-CS-045
Vrije Universiteit Amsterdam, August 2008

ABSTRACT

We introduce our 60-node indoor wireless sensor network testbed, distributed among six office rooms. We outline the hardware architecture of the testbed and discuss the basic properties of the internode connectivity graph. In particular, we show that a node in the testbed has from 8 to 31 high-quality neighbors and that the network diameter is between 4 and 5 for the lowest transmission power. The testbed is currently being used as the main evaluation platform for novel wireless sensor network protocols and systems.

Faculty of Sciences
Vrije Universiteit Amsterdam
De Boelelaan 1081A
1081 HV Amsterdam
The Netherlands

vrije Universiteit amsterdam

Contents

1 Introduction

Composed of thousands of tiny low-power embedded devices that collaboratively monitor the environment, guide vehicles, and predict potential faults in buildings, bridges, roads, and rails, wireless sensor networks (WSNs) are an important tier in the IT ecosystem [1]. They have emerged to realize the vision of an embedded Internet, in which networks formed by wirelessly interconnected computing devices provide detailed instrumentation over many points in large physical spaces. Such instrumentation constitutes a powerful tool that can potentially transform whole fields of science, engineering, and manufacturing [1].

However, due to inherent limitations of wireless sensor nodes and their interactions with the surrounding environment, sound evaluation of systems and protocols devised for WSNs is challenging. Real-world deployments to date have made it evident that in WSNs the divergence between practice and theory is even more pronounced than in other large-scale distributed systems [2, 3, 4, 5, 6]. In particular, low-power wireless communication is extremely unpredictable and, despite significant research efforts, no models that accurately capture all its intricacies exist [7, 8, 9, 10, 11]. Likewise, it is extremely difficult to accurately model the lifetime of a WSN-based system as the power consumption of a node depends on a number of dynamically changing factors, such as the number of cycles burned by the CPU, the current power mode the CPU is in, the activity on the radio channel, and the remaining battery capacity, to name a few. Such unpredictability and complexity severely limit the accuracy obtained by evaluating a WSN-based system designs analytically or with simulations. Therefore, sound evaluation that predicts the real-world system behavior relatively well should be performed with a real implementation running on real hardware.

In this technical report, we present a 60-node indoor testbed we have deployed for conducting such evaluations. Rather than comparing our testbed with the existing ones, we outline the architecture and general properties of the testbed. The objective of this report is thus describing our main experimental platform for WSN-oriented research.

The rest of the paper is structured as follows. First, in Section 2, we outline the hardware architecture and the organization of our testbed. Then, in Section 3, we show basic network properties, which we obtained through experiments. Finally, in Section 4 we conclude and discuss possible future work.

2 Hardware and Organization

The main hardware components constituting our testbed are presented in Table 1. The testbed includes 60 TelosB-class nodes [12] (see also Fig. 1 and Table 2). Ten of the nodes have been purchased at Moteiv Corporation,[1] and each of them contains a full sensor suite. The remaining fifty come from Crossbow Technology Inc.:[2] twenty of them are with and thirty are without the sensor suite.

Component	Quantity
TelosB (Moteiv) w/ sensors	10
TelosB (Crossbow) w/ sensors	20
TelosB (Crossbow) w/o sensors	30
USB hubs	16
PCs	6

Table 1. The main hardware components of our testbed.

Fig. 1. A TelosB node (from Moteiv) with the sensor suite.

The nodes are located on the fourth floor of the Faculty of Sciences of Vrije Universiteit Amsterdam. They are dispersed among six rooms as depicted in Fig. 2. Such a distribution impacts wireless internode connectivity, as discussed in the next section, thereby providing a realistic network topology. In addition, it allows for collecting environmental data from different parts of the building: southern, eastern, and western.

To avoid troublesome battery-based power supply and to enable low-overhead node retasking and statistic reporting, each of the nodes is connected to a PC using a USB network. This network consists of a number of cables and hubs connected to six Pentium III PCs, one PC per room. The wired USB network is used only as a power supply and a reliable transport backbone for protocol statistics and control commands. The protocols evaluated on our testbed, in contrast, use the standard wireless communication. In this way the interference between statistic gathering and protocol operation is minimized.

3 Basic Network Properties

To obtain basic connectivity properties of our network, we have written and deployed a simple TinyOS 2.0 application.

[1] `http://www.moteiv.com/`

[2] `http://www.xbow.com/`

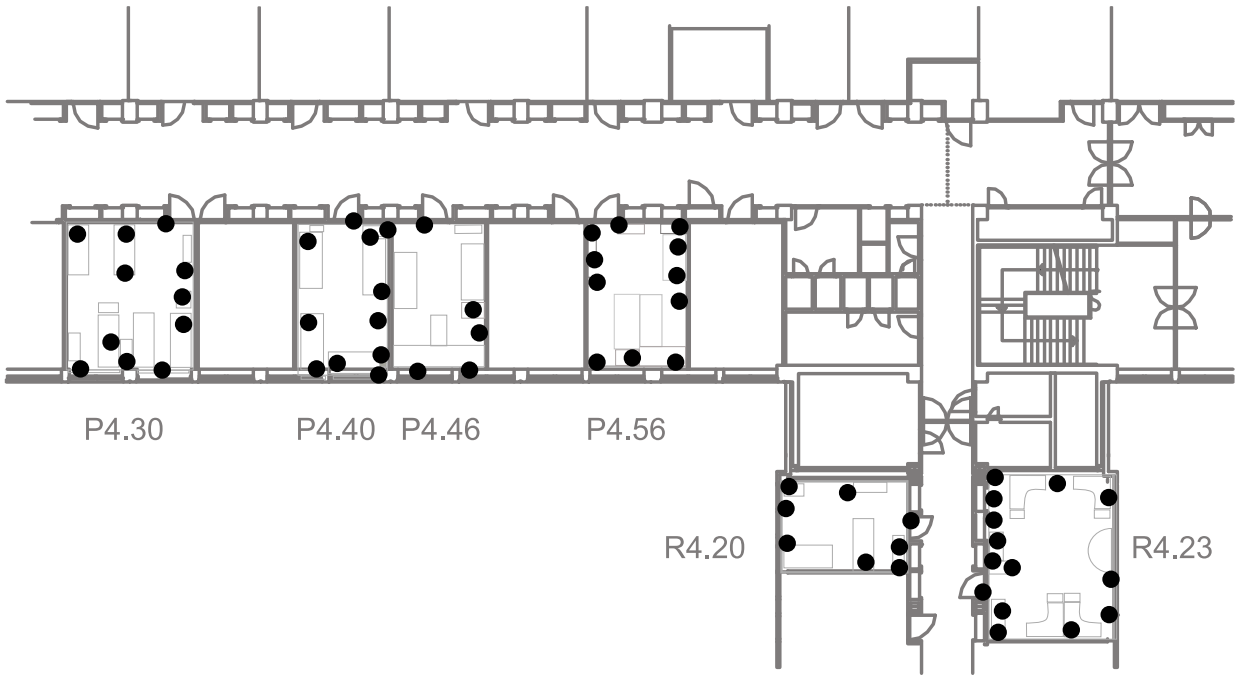

Fig. 2. The node placement in our testbed. The exact mapping of nodes to rooms is presented in Appendix A

3.1 Experimental Setup

The test application works in periods. In the default configuration, each period lasts 30 seconds and the clocks of the nodes are not synchronized. In every period, each node broadcasts a heartbeat message at a uniformly random time moment within the period. A heartbeat message of a node consists of the node's identifier (2 bytes), a sequence number (4 bytes), and a *"Hello world from Konrad Iwanicki!"* string (34 bytes including the terminating zero).

Nodes that hear the message transmission record the following metadata: the local time, the sender's identifier and sequence number, and the information about the received signal for the message, as provided by the CC2420 radio chip in the form of the received signal strength indicator (RSSI) and the link quality indicator (LQI). The recorded message metadata are placed in a metadata queue.

At the end of every period, all the records from a node's metadata queue are transmitted over the USB network to the node's corresponding PC, where they are logged by a Java frontend of our application. Therefore, ultimately the log of each node contains fine-grained data specifying all the messages the node has received and the signal quality when receiving these messages. By using such data, we are able to accurately assess wireless internode connectivity.

We have conducted four deployments of the application, each lasting at least 24 hours. In this way, we could measure connectivity during busy daytime, when we expected wireless interference caused by student and employee laptops, as well as during quiet nights, when the building was deserted (surprisingly, there was virtually no difference). Each deployment varied in the node transmission power, which was set globally for all nodes. We used the following transmission power settings: -25 dBm, -15 dBm, -5 dBm, and 0 dBm. This allowed us to investigate the impact of the transmission range on the network density and diameter. Since during the experiments the 5 nodes in room P4.46 were not deployed, we used only the remaining 55 nodes in other rooms.

3.2 Experimental Results

For each pair of nodes, in each direction, we have analyzed RSSI, LQI, and the packet loss rate (PLR), which was obtained by examining the sequence numbers of the received messages. In this way, considering for instance the average RSSI, we obtained a connectivity matrix in which a cell in row i and column j represents the average RSSI value for messages received by node i from node j. Such sample matrices for the transmission power of -25 dBm are depicted in Fig. 3, Fig. 4, and Fig. 5.

These figures show inherent clustering between the nodes. Nodes in the same room are likely to be connected with a high-quality link (high RSSI and LQI values, and a low PLR value). There are obviously some exceptions, such as nodes 02 and 03 or nodes 48 and 45. In addition, some nodes in every room can communicate with some nodes in other rooms, which ultimately results in a single connected network. Finally, all three link quality metrics are strongly correlated, again with an exception of a few links, which

Metric	Value	Remarks
MODULE		
Processor Performance	16-bit RISC	
Program Flash Memory	48K bytes	
Measurement Serial Flash	1024K bytes	
RAM	10K bytes	
Configuration EEPROM	16K bytes	
Serial Communications	UART	0-3V transmission levels
Analog to Digital Converter	12 bit ADC	8 channels, 0-3V input
Digital to Analog Converter	12 bit DAC	2 ports
Other Interfaces	Digital I/O,I2C,SPI	
Current Draw	1.8 mA 5.1 μA	Active mode Sleep mode
RF TRANSCEIVER		
Frequency band*	2400-2483.5 MHz	ISM band
Transmit (TX) data rate	250 kbps	
RF power	-25 dBm to 0 dBm	
Receive Sensitivity	-90/-94 dBm	minimal/typical
Outdoor Range	75 m to 100 m	Inverted-F antenna
Indoor Range	20 m to 30 m	Inverted-F antenna
Current Draw	23 mA 21 mA 1 μA	Receive/Listen Transmit (at 0 dBm) Off
SENSORS (Optional)		
Visible Light Sensor Range	320 nm to 730 nm	Hamamatsu S1087
Visible to IR Sensor Range	320 nm to 1100nm	Hamamatsu S1087-01
Humidity Sensor Range Resolution Accuracy	0-100% RH 0.03% RH $\pm$ 3.5%	Sensirion SHT11 RH Absolute RH
Temperature Sensor Range Resolution Accuracy	-40°C to 123.8°C 0.01°C $\pm$ 0.5°C	Sensirion SHT11 at 25°C
MECHANICAL		
Battery	2X AA batteries	Attached pack
User Interface	USB	v1.1 or higher
Size (in) (mm)	2.55 x 1.24 x 0.24 65 x 31 x 6	w/o battery pack w/o battery pack
Weight (oz) (grams)	0.8 23	w/o batteries w/o batteries

*Programmable in 1-MHZ steps, 5-MHz steps for compliance with IEEE 802.15.4/D18-2003.

Table 2. The specification of TelosB (source: Crossbow).

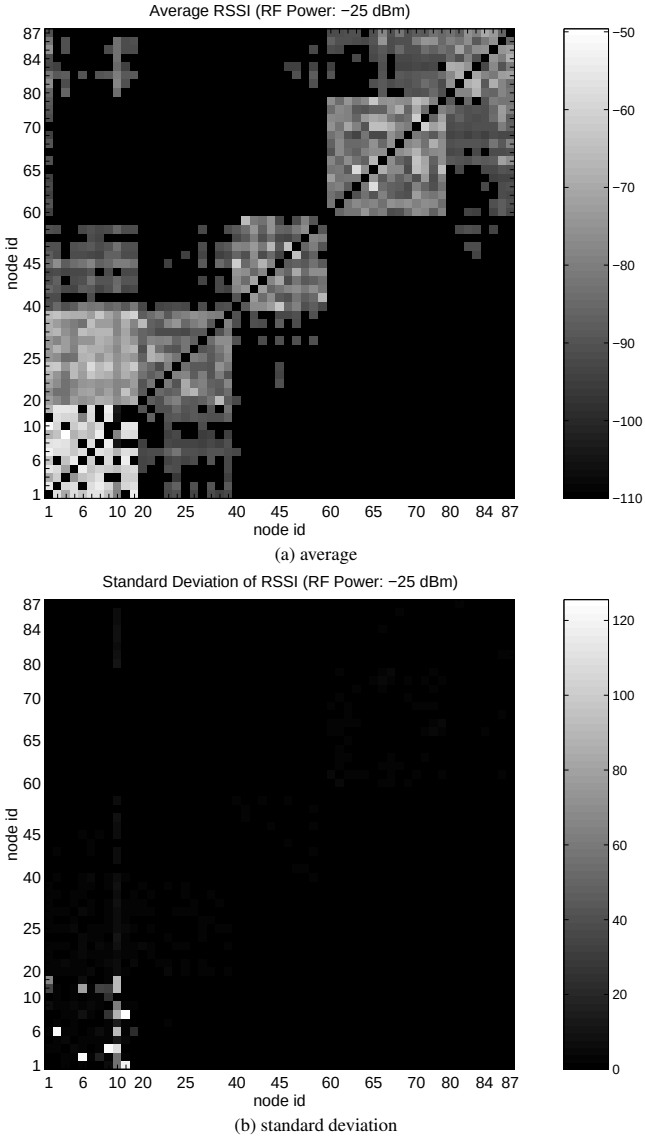

(a) average

(b) standard deviation

Fig. 3. RSSI matrices for the RF power of −25 dBm.

matches the results reported by other groups [13, 7, 9].

Based on the connectivity matrices, we also generated connectivity graphs. Each vertex in the graph corresponds to a node. There exists an edge between two vertices if the communication links between the nodes associated with these vertices meet certain requirements. For the RSSI metric, for instance, the RSSI value of the links in both directions must be at least −90 dBm. For the LQI metric, in turn, the LQI value of the links in both directions must be at least 95. Finally, for PLR metric, the packet loss rate of the links must be below 15%. These thresholds represent high-quality links and were obtained from the results reported by other groups [13, 7, 9].

For different transmission power levels and link quality metrics, we computed certain graph-theoretical properties of the connectivity graphs. The properties include: the graph diameter (the maximal shortest path between all pairs of nodes), the network density (the degree of a node in the connectivity graph), and the clustering coefficient (the ratio of the actual number of links between a node's neighbors to all the possible links between the node's neighbors). The results for different metrics and power level are shown in Table 3.

It can be observed that in all configurations the network is multi-hop with the diameter of 3 for the highest transmission power and the diameter of to 4 or 5, depending on the link quality metric, for the lowest transmission power. A node has on average more than 15 and less than 30 high-quality neighbors. Moreover, the connectivity graph is highly clustered, as mentioned previously. High clustering is an inherent feature of wireless networks as, due to the limited radio range, two nodes that have a common neighbor are very likely to be each other's neighbors as well [14].

To complete the picture, in Fig. 6 and Fig. 7 we show the distribution of node degree and path length for the PLR metric and the lowest transmission power.

The degree distribution (see Fig. 6) demonstrates that, more than 50% of the nodes have at most 17 good quality neighbors. Yet, there exist some nodes that have more than

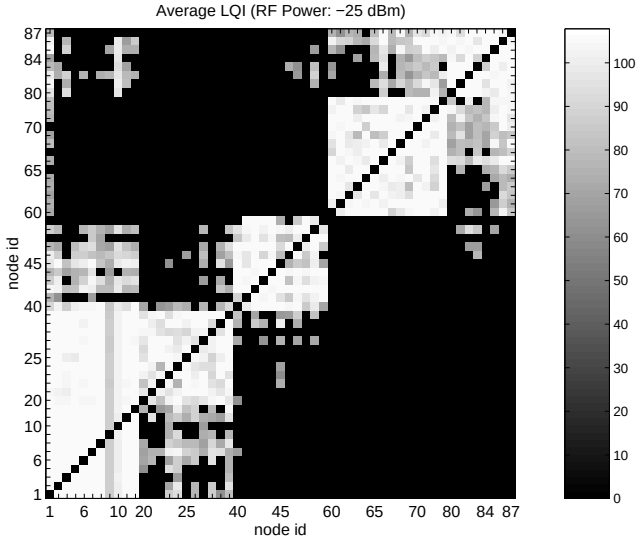

(a) average

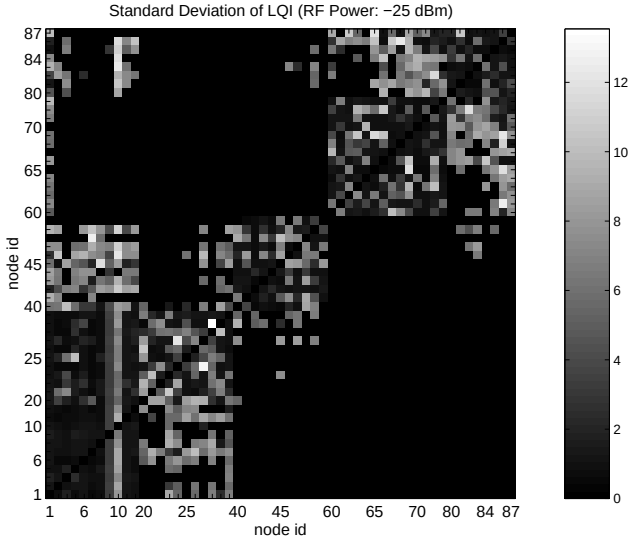

(b) standard deviation

Fig. 4. LQI matrices for the RF power of −25 dBm.

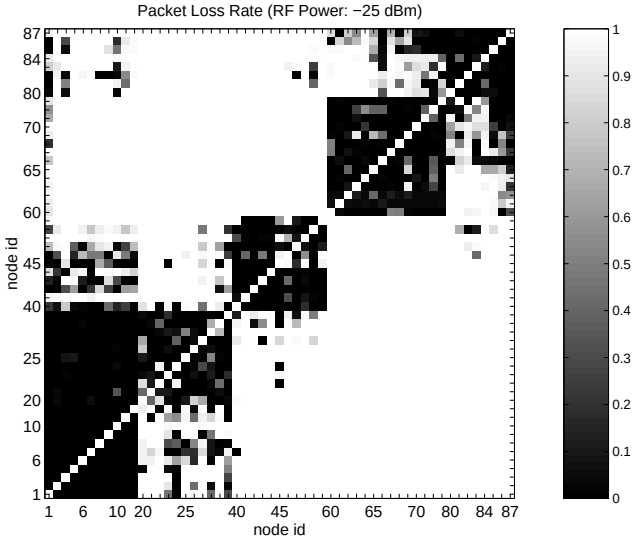

Fig. 5. The PLR matrix for the RF power of −25 dBm.

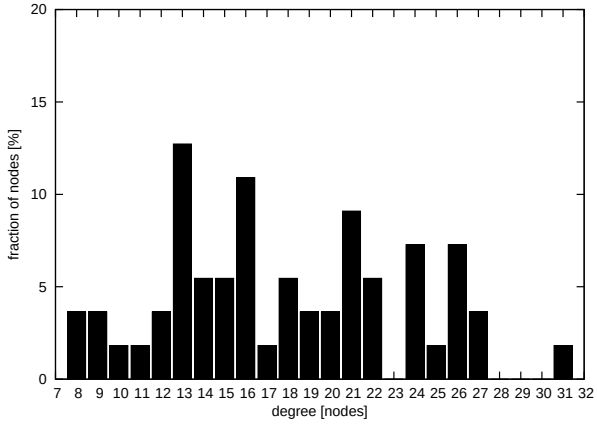

Fig. 6. Node degree distribution (RF power of −25 dBm).

25 neighbors. In general, the node degree distribution is highly non-uniform and varies between 8 and 31 neighbors. Such non-uniformity is very common in real-world WSN deployments.

The distribution of the path length (see Fig. 7) is more predictable. A bit more than 30% of all 2970 paths are one hop. This is a direct consequence of the high clustering coefficient. More than 60% of the paths are at most two hops, and nearly 90% are at most three hops. However, a relatively large fraction of paths has the same length as the network diameter. This implies that, to reach each other, many nodes must forward messages over the distance equal to the network diameter. Such long paths facilitate testing of various routing algorithms.

4 Conclusions and Future Work

We introduced our 60-node indoor WSN testbed. We outlined the hardware architecture of the testbed and presented the basic properties of the internode connectivity graph. The graph has highly non-uniform node density and many paths reaching the network diameter of 4 to 5 hops. We believe that these properties of our testbed will enable sound evaluation of many of the WSN protocols and systems we have devised or hope to devise in the future.

During our experiments, we also identified a need for automating testbed reprogramming and experiment scheduling. To this end, we plan to augment our testbed with some job scheduling software, similar to the software for cluster computers. This would allow multiple users to concurrently access the testbed, which is simply a necessity if the user population grows beyond a few persons.

RF Power	Diameter	Density		Clust. Coeff.	
		AVG	STDEV	AVG	STDEV
−25 dBm	4	19.65	5.96	0.82	0.13
−15 dBm	3	27.62	5.58	0.84	0.12
−5 dBm	3	29.22	6.07	0.83	0.12
0 dBm	3	29.87	6.85	0.80	0.11

(a) RSSI metric

RF Power	Diameter	Density		Clust. Coeff.	
		AVG	STDEV	AVG	STDEV
−25 dBm	5	17.04	5.83	0.83	0.13
−15 dBm	3	25.96	4.29	0.86	0.12
−5 dBm	3	27.51	4.69	0.86	0.12
0 dBm	3	28.31	5.72	0.86	0.12

(b) LQI metric

RF Power	Diameter	Density		Clust. Coeff.	
		AVG	STDEV	AVG	STDEV
−25 dBm	4	17.87	5.57	0.78	0.13
−15 dBm	3	27.00	4.78	0.84	0.12
−5 dBm	3	28.24	4.75	0.84	0.11
0 dBm	3	29.04	5.03	0.82	0.11

(c) PLR metric

Table 3. Connectivity graph properties depending on RF power.

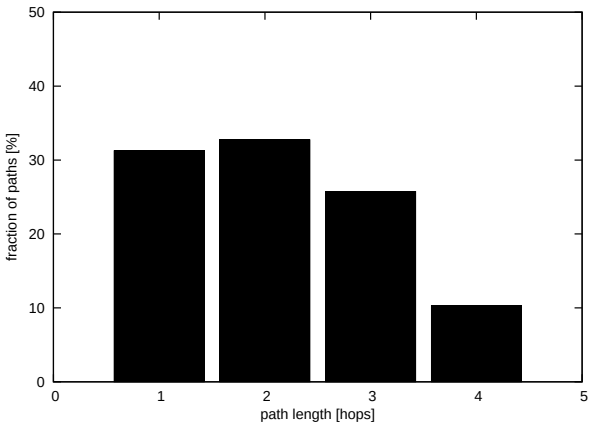

Fig. 7. Path length distribution (RF power of −25 dBm).

Acknowledgments

We would like to thank all people that allowed us to deploy the sensor nodes in their offices: A. Bakker, P. Costa, J. Domaschka, D. Gavidia, R. Kemp, K. Mitrokotsa, J. Napper, N. Paul, V. Rai, L. Ranaweera, M. Rieback, and J. Słowińska. Without their cooperation the deployment could not have been completed.

References

[1] D. Culler and W. Hong, Eds., *Communications of the ACM*, June 2004, vol. 47, no. 6, ch. Wireless Sensor Networks, pp. 30–57.

[2] R. Szewczyk, A. Mainwaring, J. Polastre, J. Anderson, and D. Culler, "An analysis of a large scale habitat monitoring application," in *Proceedings of the Second ACM Int. Conf. on Embedded Networked Sensor Systems (SenSys)*, Baltimore, MD, USA, November 2004, pp. 214–226.

[3] G. Tolle and D. Culler, "Design of an application-cooperative management system for wireless sensor networks," in *Proceedings of the Second European Workshop on Wireless Sensor Networks (EWSN)*, Istanbul, Turkey, January 2005.

[4] N. Ramanathan, E. Kohler, and D. Estrin, "Towards a debugging system for sensor networks," *International Journal of Network Management*, vol. 15, no. 4, pp. 223–234, July 2005.

[5] K. Langendoen, A. Baggio, and O. Visser, "Murphy loves potatoes: Experiences from a pilot sensor network deployment in precision agriculture," in *Proceedings of the Twentieth IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Rhodes Island, Greece, April 2006.

[6] G. Werner-Allen, K. Lorincz, J. Johnson, J. Lees, and M. Welsh, "Fidelity and yield in a volcano monitoring sensor network," in *Proceedings of the Seventh USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, Seattle, WA, USA, November 2006, pp. 381–396.

[7] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis, "Some implications of low-power wireless to IP routing," in *Proceedings of the Fifth ACM Workshop on Hot Topics in Networks (HotNets)*, Irvine, CA, USA, November 2006, pp. 31–36.

[8] M. M. Holland, R. G. Aures, and W. B. Heinzelman, "Experimental investigation of radio performance in wireless sensor networks," in *Proceedings of the Second IEEE Workshop on Wireless Mesh Networks (WiMesh)*, Reston, VA, USA, September 2006, pp. 140–150.

[9] K. Srinivasan and P. Levis, "Rssi is under appreciated," in *Proceedings of the Third ACM Workshop on Embedded Networked Sensors (EmNets)*, Cambridge, MA, USA, May 2006.

[10] H. Lee, A. Cerpa, and P. Levis, "Improving wireless simulation through noise modeling," in *Proceedings of the Sixth International Conference on Information Processing in Sensor Networks (IPSN)*, Cambridge, MA, USA, April 2007.

[11] R. Fonseca, O. Gnawali, K. Jamieson, and P. Levis, "Four-bit wireless link estimation," in *Proceedings of the Sixth ACM Workshop on Hot Topics in Networks (HotNets)*, Atlanta, GA, USA, November 2007.

[12] J. Polastre, R. Szewczyk, and D. Culler, "Telos: Enabling ultra-low power wireless research," in *Proceedings of the Fourth International Conference on Information Processing in Sensor Networks (IPSN) (Poster Session)*, Los Angeles, CA, USA, April 2005, p. 48.

[13] A. Woo, T. Tong, and D. Culler, "Taming the underlying challenges of reliable multihop routing in sensor networks," in *Proceedings of the First ACM Int. Conf. on Embedded Networked Sensor Systems (SenSys)*, Los Angeles, CA, USA, November 2003, pp. 14–27.

[14] S.-Y. Ni, Y.-C. Tseng, Y.-S. Chen, and J.-P. Sheu, "The broadcast storm problem in a mobile ad hoc network," in *Proceedings of the Fifth ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, Seattle, WA, USA, August 1999, pp. 151–162.

A Node to Room Mapping

Figure 8-Fig. 12 present the exact mapping of node identifiers to physical node locations. This data may be useful when analyzing connectivity matrices.

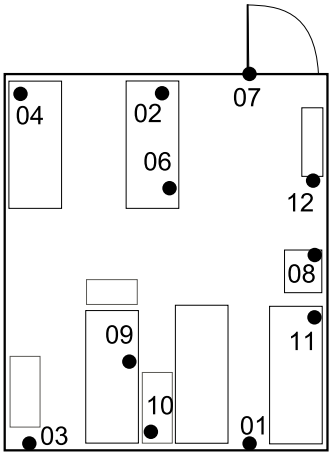

Fig. 8. Room P4.30.

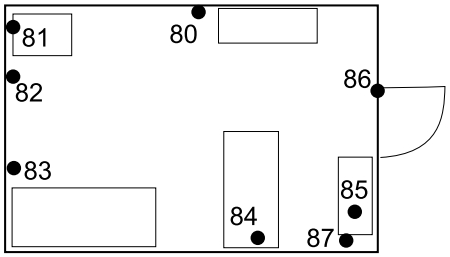

Fig. 11. Room R4.20.

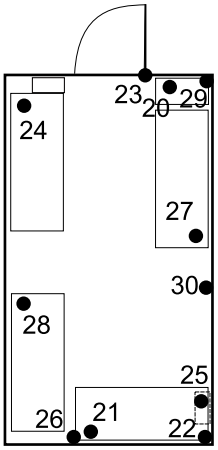

Fig. 9. Room P4.40.

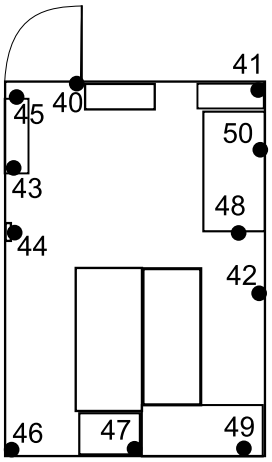

Fig. 10. Room P4.56.

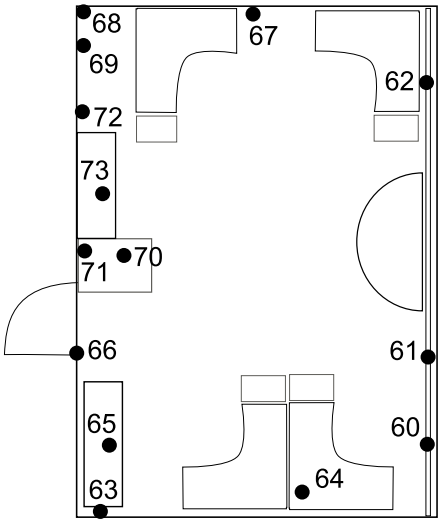

Fig. 12. Room R4.23.