Data Collection in Wireless Sensor Networks for Noise Pollution Monitoring

Luca Filipponi[1], Silvia Santini[2] *, and Andrea Vitaletti[1] **

[1] Dipartimento di Informatica e Sistemistica "A. Ruberti"
SAPIENZA Università di Roma, Rome, Italy
{filipponi,vitale}@dis.uniroma1.it
[2] Institute for Pervasive Computing
ETH Zurich, 8092 Zurich, Switzerland
santini@inf.ethz.ch

Abstract. Focusing on the assessment of environmental noise pollution in urban areas, we provide qualitative considerations and experimental results to show the feasibility of wireless sensor networks to be used in this context. To select the most suitable data collection protocol for the specific noise monitoring application scenario, we evaluated the energy consumption performances of the CTP (Collection Tree Protocol) and DMAC protocols. Our results show that CTP, if used enabling the LPL (Low Power Listening) option, provides the better performances trade-off for noise monitoring applications.

1 Environmental noise monitoring

Conservative estimations give in about 300 millions the number of citizens within the European Community that are exposed to alarming levels of noise pollution [1]. Raising the public's awareness of this problem, the Directive 2002/49/EC of the European Parliament has made the avoidance, prevention, and reduction of environmental noise a prime issue in European policy. To better assess the extension of the problem, the European Commission required member states to regularly provide an accurate mapping of environmental noise levels for all urban areas with more than 250′000 inhabitants. While current noise maps are mostly based on sparse data and ad-hoc noise propagation models, a recent position paper by the Commission has stressed that *"every effort should be made to obtain accurate real data on noise sources,"* [2, p.6]. The demand for accurate data about noise exposure levels will increase dramatically, as this statement makes its way into mandatory regulation. Nowadays noise measurements in urban areas are mainly carried out by designated officers that collect data at a location of interest for successive analysis and storage, using a sound level meter or similar device. This manual collection method does not scale as the demand for higher

* Partially supported by the Swiss National Science Foundation (NCCR-MICS, grant number 5005-67322).
** Partially supported by the FRONTS FET Project (215270).

granularity of noise measurements in both time and space increases. Instead, a network of cheap wireless sensor nodes deployed over the area of interest could collect noise pollution data over long periods of time and autonomously report it to a central server through the sensor's on-board radio, requiring human intervention only to install and possibly subsequently remove the sensing devices. Collected noise data is typically stored in a land register and used, together with additional information about existing noise sources, to feed computational models that provide extrapolated noise exposure levels for those areas for which real data is unavailable. Even if this assessment procedure is still compliant with European regulations, today's computational models often fail to provide accurate estimations of the real noise pollution levels[3]. Indeed, while the free propagation properties of noise generated from typical noise sources are well understood [3], shadowing and reflection effects hinder accurate estimation of noise levels in complex urban settings. The accuracy of computed noise levels could be easily verified and improved by installing a wireless sensor network at those locations for which computational models are likely to provide inaccurate estimations.

2 Requirements

Before going into further details we would like to summarize the main requirements a wireless sensor network must comply with to be used for noise pollution monitoring applications.

Hardware. The high sampling rate required to properly capture acoustic signals ($\sim$ 32kHz) appears prohibitive for resource poor sensor nodes. However, commercially available platforms are able to support the required sampling rate, as long as scheduling with radio communication is properly managed. Nevertheless, to overcome this problem the most suitable solution consists in delegating sampling and signal processing to dedicated hardware and let the actual sensor node only deal with communication and possibly optimization of data collection.

Sampling. Noise levels are time-weighted averages of acoustic power, computed over variable time intervals. Sampling of noise levels may occur at sampling rates varying from fractions up to multiples of 1 Hz. For the preparation of noise maps, a spacing of about 3 meters between sampling points (i.e. sensor nodes) is recommended [2]. Observe that this results in a quite dense network deployed on a relatively small area (e.g., the internal façade of a building).

Data rate and latency. Collection of noise data does not require sensor readings to be immediately reported to the sink. Therefore, latency in packet delivery is a secondary issue for the optimization of network performances. In typical scenarios, sensor nodes generate 1 value/sec and a single packet can convey several noise samples by means of aggregation.

Network lifetime. To cover typical variability patterns of noise levels, measurements should ideally extend for few weeks.

[3] The authors are indebted to Hans Huber and Fridolin Keller of the department for environmental noise protection of the city of Zurich, who pointed this out in a personal in-depth interview.

Network topology. For the purpose of noise mapping, the assessment points used to measure noise levels (thus, the physical topology of the network) shouldn't change during data collection. Nevertheless, due to the spatial distribution of the nodes, reporting data to a central sink may require multi-hop communication.
Synchronization. Noise readings collected by different nodes must be ordered over a global timescale for proper processing and visualization. Since the specific network topology and data collection protocol may introduce a variable and unbounded latency on data delivery, an adequate synchronization mechanism should be adopted to allow for a correct time ordering.

3 Data collection

To select an adequate platform to collect noise levels data, we tested the feasibility of three different hardware solutions. We considered the *Tmote Sky* prototyping platform from Moteiv Corp. equipped with either the *EasySen SBT80* multi-modality sensor board or with a custom-made noise level meter. Furthermore, we experimented with the *Tmote invent* platform, also from Moteiv Corp., which features an on-board microphone, as well as a powerful signal conditioning circuitry [4]. Taking care of sampling and processing the captured acoustic signal, the customized noise level meter behaves as an external sensor able to output noise level readings expressed in dB (with total nominal error less than 3 dB). Since the noise level is actually a time-weighted average of the captured acoustic power, its value can be sampled at much lower frequency (e.g., 1 Hz) than the acoustic signal itself. The use of a customized noise level meter allows therefore to remove the burden of computational and energy expensive operations from the sensor node itself, and represents our preferred solution at this first prototyping stage.

After being captured, noise level data needs to be transmitted to a central sink for permanent storage and further processing, imposing a typical convergecast pattern on network communication. While the design of energy-efficient medium access control received considerable attention within the wireless sensor networks research community [5], only few collection layer implementations are currently available, these including MintRoute [6] and CTP (Collection Tree Protocol) [7]. In particular, CTP is an implementation of BMAC [8] with an optional low power listening (LPL) option on trees. LPL is a power saving technique that allows to move the major costs of radio communication from receivers to transmitters by avoiding idle listening, which is known to be the main source of energy wasting in wireless sensor networks communication. To understand the performances of different data collection protocols in our specific application scenario, we evaluated both the CTP and the DMAC convergecast protocol [9]. We consider the comparison of these two protocols particularly interesting since, at the media access level, they respectively implement a contention-based and TDMA-based (Time Division Multiple Access) approach, which are the mostly exploited MAC techniques in wireless sensor networks. We embedded both implementations of the data collection layer in our software prototype and analyzed their performances in terms of energy consumption. While an implementation of CTP is available in the tinyOS-2.x repository [7], we implemented the DMAC protocol on our own.

4 Assessment and analysis of protocols' performances

To analyze the energy consumption of the two data collection protocols under consideration, we connected a Tmote sensor node to the Rhode & Schwarz dual-channel analyzer/power supply NGMO2,an instrument that can accurately measure the current drain associated with all the states of a node. We then measured the node current drain during transmission and reception states, for both the CTP and DMAC protocols, using the same experimental set-up exploited in [10]. This simple setting, a single-hop network made of two nodes, allowed us to gain a clear understanding of all the major energetic aspects involved in nodes communications and to determine upper and lower bounds on nodes energy consumption. A more realistic experimental setting, considering multi-hop topologies and the effect of collisions, will be considered in further investigations. We would like to point out that measurements of power consumption in wireless sensor networks typically rely on indirect measurement methods, such as counting the number of transmitted packets or CPU duty cycles, which provide limited accuracy as pointed out in [11]. Our work aligns with few other examples in providing direct power consumption measurements [10–12]. Figure 1(a) shows the current consumption of a node running the CTP_{NoLPL} protocol, characterized by the transmissions spikes, occurring every second. While in idle listening, the node drains about 19mA, which represents a clear energy waste considering that the total average current drain is about 19.5 mA. Enabling the LPL option allows to dramatically reduce the power consumption of the CTP protocol, as shown in figure 1(c). The current drain while the radio is in sleep mode is negligible. Every 250ms (the value we set for nodes sleep period), the node wakes up the radio and samples the channel, as foreseen by the LPL mechanism. If it overhears a transmission on the medium, it keeps the radio active until reception is successfully completed, otherwise switches-off the radio immediately. The spikes in figure 1(c) are associated to the sampling activity, while one complete reception cycle is clearly visible in the first segment of the plot. The length of the transmission phase is considerably longer than the one observed for CTP_{NoLPL}, since LPL moves the cost of communication from the receiver to the transmitter. Nevertheless, under the same traffic load of one packet per second, the total average current drain of the CTP_{LPL} protocol is only 5mA, which represents an energy saving of 75% with respect to the 19.5mA spent on average by the CTP_{NoLPL}. The CTP_{LPL} protocol can achieve even lower average current drains using longer sleep intervals that, however, will also increase packet latency. For both CTP_{NoLPL} and CTP_{LPL}, the measured current drain values represents lower bounds on the protocol's energy consumption, since in the considered experimental setting no collisions and packet forwarding occur.

For the DMAC protocol, the lower bound on energy consumption simply corresponds to the energy required to transmit (receive) a single packet per

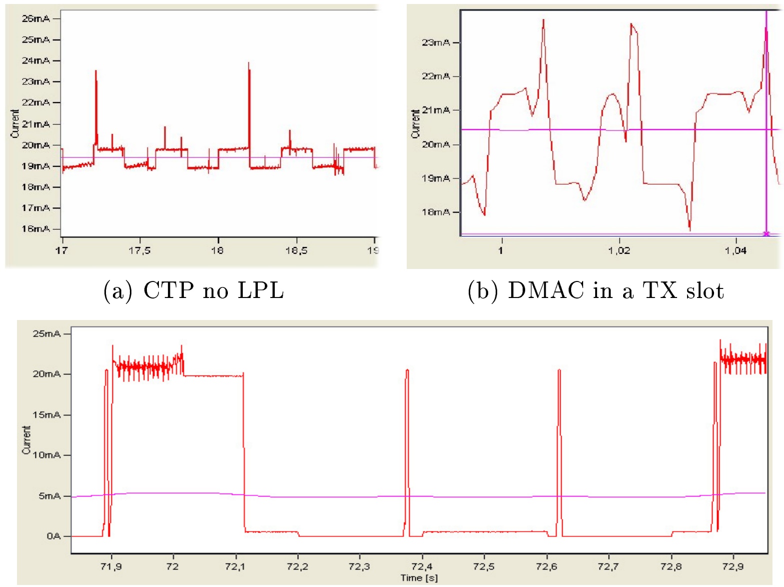

(a) CTP no LPL (b) DMAC in a TX slot

(c) CTP with LPL (sleep interval 250ms)

Fig. 1. Current drain

period. Determining the upper bound is less trivial and requires some additional considerations. Nodes running the DMAC protocol can, at each time instants, be in sleep, receive or transmit state. The power consumption in sleep state is negligible, so we excerpt it from the current analysis. To provide an upper bound on DMAC's power consumption, we forced nodes in transmission state to transmit for an entire period, and consequently nodes in receive state to remain active for the duration of such a period. Figure 1(b) shows the current drain of the node in transmission state, in which the average current consumption is about 20.5mA (about 21.5mA during reception). A correct estimate of the average energy consumption, requires to take into consideration the time spent in sleep state. If we consider a 10-layers network, with communication periods of 100ms, the DMAC's *time-weighted* average energy consumption is approximatively the same as for CTP_{LPL}.

5 Conclusions

Wireless sensor networks can provide a cheap and flexible infrastructure to support the collection of fine-grained noise pollution data, which is essential for the preparation of noise maps and for the validation of noise pollution models. Besides testing commercially available sensor platforms, we designed and developed a customized noise level meter that allows us to delegate costly noise levels computations to dedicated hardware. We then selected two data collection protocols, CTP (considered both with and without the LPL option) and DMAC, and performed direct measurements of their energy consumption using a simple, though representative, network topology. The CTP_{NoLPL} protocol exhibits the highest energy consumption (and lowest latency), while enabling the LPL option allows to save about 75% of the total energy (though at the cost of increased latency). DMAC's observed performances are comparable to those of CTP_{LPL}, though they depend on the number of tree levels the protocol builds up for routing. For our specific application scenario we eventually selected the simple CTP_{LPL} as the most suitable data collection protocol.

References

1. European Commission Green Paper: Future noise policy. Com (96) 540 final (November 1996)
2. European Commission Working Group Assessment of Exposure to Noise (WG-AEN): Good practice guide for strategic noise mapping and the production of associated data on noise exposure (January 2006)
3. Bies, D.A., Hansen, C.H.: Engineering Noise Control: Theory and Practice. 3rd edn. Spon Press (Taylor &Francis Group), London and New York (2003)
4. Santini, S., Ostermaier, B., Vitaletti, A.: First experiences using wireless sensor networks for noise pollution monitoring. In: Proceedings of the 3rd ACM Workshop on Real-World Wireless Sensor Networks (REALWSN'08), Glasgow, United Kingdom, ACM (April 2008)
5. Demirkol, I., Ersoy, C., Alagöz, F.: Mac protocols for wireless sensor networks: a survey. IEEE Communications Magazine **44**(4) (April 2006) 115–121
6. Woo, A., Tong, T., Culler, D.: Taming the underlying challenges of reliable multihop routing in sensor networks. In: Proceedings of the 1st international conference on Embedded networked sensor systems (SenSys'03), New York, USA, ACM (2003)
7. Fonseca, R., Gnawali, O., Jamieson, K., Kim, S., Levis, P., Woo, A.: Tinyos enhancement proposals (tep) 123: The collection tree protocol (ctp)
8. Polastre, J., Hill, J., Culler, D.: Versatile low power media access for wireless sensor networks. In: Proceedings of the 2nd international conference on Embedded networked sensor systems (SenSys'04), New York, NY, USA, ACM (2004) 95–107
9. Lu, G., Krishnamachari, B., Raghavendra, C.S.: An adaptive energy-efficient and low-latency mac for data gathering in wireless sensor networks. In: Int. Workshop on Algorithms for Wireless, Mobile, Ad Hoc and Sensor Networks (WMAN), Santa Fe, NM, USA (April 2004)
10. van Dam, T., Langendoen, K.: An adaptive energy-efficient mac protocol for wireless sensor networks. In: Proceedings of the 1st international conference on Embedded networked sensor systems (SenSys'03), New York, USA (2003) 171–180
11. Landsiedel, O., Wehrle, K., Götz, S.: Accurate prediction of power consumption in sensor networks. In: Proceedings of the Second IEEE Workshop on Embedded Networked Sensors (EmNetS-II), Sidney, Australia (May 2005)
12. Hohlt, B., Doherty, L., Brewer, E.: Flexible power scheduling for sensor networks. In: Proceedings of the third international symposium on Information processing in sensor networks (IPSN '04), New York, NY, USA, ACM (2004) 205–214