Adaptive Sampling for Energy Conservation in Wireless Sensor Networks for Snow Monitoring Applications

Cesare Alippi[*], Giuseppe Anastasi[§], Cristian Galperti[*], Francesca Mancini[§], Manuel Roveri[*]

[*]*Dip. di Elettronica e Informazione*
Politecnico di Milano, Italy
{lastname}@elet.polimi.it

[§]*Dept. of Information Engineering*
University of Pisa, Italy
{firstname.lastname}@iet.unipi.it

Abstract

Energy conservation techniques for sensor networks typically rely on the assumption that data sensing and processing consume considerable less energy than communication. This assumption does not hold in some practical application scenarios, where ad hoc developed sensor units require power consumption comparable with, or even larger than, that of the radio. In this paper we focus on an embedded sensor for monitoring snow composition in mountain slopes for avalanche forecasting. To lower the sensor energy consumption we propose an adaptive sampling algorithm able to dynamically estimate the optimal sampling frequency of the signal to be monitored. In turn, this minimizes the activity of both the sensor and the radio (hence saving energy) while maintaining an acceptable accuracy on the acquired data. Simulation experiments show that the suggested solution can save up to 97% of the energy consumed for sensing when the sensor is always on, while maintaining the error at acceptable levels.

1. Introduction

A sensor network typically consists of a large number of sensor nodes deployed over a geographical area. Each node is a low-power device that integrates processing, sensing and wireless communication abilities. Sensor nodes acquire information from the surrounding environment, process locally the data, and/or send them to one or more collection points (base stations) [1]. A main issue for a credible deployment of sensor units is energy consumption as sensor nodes are, generally, battery powered. The battery has limited capacity and, often, cannot be replaced or recharged, due to environmental or cost constraints. Therefore, the design of any component in the sensor network should address minimization of the energy consumption. As the radio component of sensor nodes mainly account for energy consumption (in addition to flash memory storage), several energy management techniques have been proposed that aim at minimizing the radio activity. They include data compression [2][3] and aggregation [4][5], predictive monitoring [6], topology management [7][8], adaptive duty cycle [9], just to name the few. All the above techniques, however, rely on the assumption that the energy required for data sensing and processing is negligible with respect to the energy consumed in data communication. Despite the fact that this assumption holds in most applications, some others require ad-hoc sensors with a power consumption comparable with, or even larger than that of the radio [10]. In this paper we focus on a sensor board for snow monitoring applications; data, suitably aggregated, can be used to forecast possible avalanches. Energy savings mechanisms for the radio have been extensively studied in the past (e.g., see [11]). Here, we approach the problem at the sensor board level. Even if we focus the attention to snow sensor to make the presentation easy to follow, the methodology is quite general and can be applied – possibly with minor adaptation- to any unit characterized by sensors with non negligible energy consumption. Energy conservation at the sensor level can be achieved by using an adaptive duty cycle approach, which consists in (i) switching off the sensor board between two consecutive samples, and (ii) using the optimal sampling frequency for the physical quantity to be monitored. Specifically, we propose an Adaptive Sampling Algorithm for estimating the optimal sample frequency. The basic idea is to find dynamically the minimum sample rate compatible with the monitored signal. By reducing the sampling rate, the algorithm also reduces the amount of data to be transmitted and, hence, the energy consumed by the radio. The concept of adaptive sampling is not new [12][13][14]. However, previous proposals are aimed

Work supported by the Italian Ministry for Education and Scientific Research (MIUR) under the FIRB ArtDeco and PRIN WiseMaP projects.

1-4244-1455-5/07/$25.00 ©2007 IEEE

mainly at optimizing the communication between sensor nodes and Base Station (BS). In [12] the sampling rate is adapted to the characteristics of the data stream so that more bandwidth is allocated to sensor nodes with larger activity. In [13] adaptive sampling is used to support routing. However, no specific sampling algorithm is proposed. In [14] adaptive sampling basically consists in activating the appropriate number of sensor nodes to achieve a target error level, depending on spatial correlation and activity. Instead, our proposal relies on the CUSUM change detection test [15] and is aimed at reducing energy consumption for both sensing and communication. We show by a preliminary analysis that it is practical, and actually it reduces the energy consumption of the snow sensor up to 97%.

The paper is organized as follows. Section 2 describes the snow sensor board considered in this paper. Section 3 introduces the Adaptive Sampling Algorithm. Section 4 describes the simulation environment used for performance analysis, while Section 5 presents simulation results.

2. Snow Monitoring Applications

Monitoring the status of snow coverage is an important public protection issue which allows experts for forecasting avalanches. Moreover, information regarding the snow coverage is welcome to quantify the potential presence of water to be subsequently released and used; such information is fruitful for optimal planning hydropower generation. To gather information related to snow instabilities on mountains slopes it is crucial to identify the composition of layers of snow at different heights from the ground. Snow is a mix of ice, water and air, and its dielectric constant, measured at different frequencies in the range [0.1, 100] kHz, varies with the percentage of water and ice in the mix. Therefore, by estimating the snow dielectric constant over time, it is possible to achieve information about the composition of different snow layers [16]. The snow sensor is thus a capacitance reading unit composed of an ad-hoc engineered probe to be left on the mountain (for example fixed on a pole) measuring snow capacitance, and an electronic injection board capable of driving the probe and measuring capacity at different frequencies of excitation (Patent pending No. US 2006/0192568 A1).

For each sample cycle the sensor provides measurements of snow capacitance at 100 Hz (low frequency) and 100 kHz (high frequency). At the same time a second sensor provides a measurement of the ambient temperature. The three readings are then passed to the sensor node to be packed in a single message and sent over the wireless channel. For each measurement, the injection board electronics of the snow sensor makes several procedures (calibration, electrode pre-charging, charge sharing) in a cyclic way to obtain a reasonably stable and reliable measure. This activity makes the sensor very energy consuming: for instance, by sampling data every 15s, the average energy consumed is 880 mJ/sample. Such a high value can be explained as follows: a) the sensor is an ad-hoc sensor not optimized for energy consumption; b) the sensor is always active (no energy management is currently available on the sensor). We discovered that a good duty cycle for the sensor is around 2s for a 150 mJ/sample energy consumption. Immediately, by integrating the duty cycle concept, the energy consumption decreases by 83%.

3. Adaptive Sampling Algorithm

In this section we describe the algorithm for estimating the optimal sampling rate of snow capacitance. Nyquist [17] defines the minimum sampling frequency F_N guaranteeing the reconstruction of the sampled signal:

$$F_N > 2\ F_{max}$$

where F_{max} is the maximum frequency in the power spectrum of the signal. Unfortunately, the maximum frequency F_{max} is generally not a priori available. Moreover, in a non-stationary process, the frequency spectrum of the signal (and as a consequence the maximum frequency) may change over time. The aim of the proposed Adaptive Sampling Algorithm is thus to track the dynamic of the physical process under monitoring by adapting the sampling frequency of the sensor to the process dynamic. The proposed solution is based on the non-stationarity change detection CUSUM test.

Change detection techniques are statistical tests that assess the stationary hypothesis for a process under investigation. Here, we modified the traditional CUSUM test to assess the non-stationarity (and hence the change) of the maximum frequency of the power spectrum of the signal (and not the signal itself).

The algorithm initially estimates, through a Fast Fourier Transform (FFT), the maximum frequency $\overline{F}_{\max}$ of the signal by relying on the first W samples coming from the process (e.g., $W = 400$ samples). The estimated $\overline{F}_{\max}$ becomes the starting reference value to be contrasted for new estimates, possibly associated with changes in the maximum frequency of the signal

spectrum. Directly from the Nyquist theorem, a suitable sampling frequency for $\overline{F}_{\max}$ is

$$F_c = c\,\overline{F}_{\max} \quad \text{with} \quad c > 2$$

where c is a confidence parameter. We defined c = 2.1 for allowing the algorithm for detecting frequencies higher than $2\overline{F}_{\max}$.

To track the process dynamic, the proposed algorithm defines the two alternative hypotheses for the maximum frequency as:

$$F_{up} = \left(1 + \frac{c-2}{4}\right)\overline{F}_{\max}; \quad F_{down} = \left(1 - \frac{c-2}{4}\right)\overline{F}_{\max}.$$

During the operational life, the algorithm computes the current maximum frequency F_{curr} of the signal on sequences of W samples and the change detection test can then be applied: *if* F_{curr} is closer to F_{up} or F_{down} than $\overline{F}_{\max}$ for h consecutive samples, a change is detected in the maximum frequency of the signal and a new sampling frequency is defined. More specifically, the change detection test is composed by the following detection rule:

if $\left(\left|F_{curr} - F_{up}\right| < \left|F_{curr} - \overline{F}_{\max}\right|\right)$ *for* h *consecutive samples or if* $\left(\left|F_{curr} - F_{up}\right| < \left|F_{curr} - \overline{F}_{\max}\right|\right)$ *for* h *consecutive samples, update the sampling frequency:* $F_c = c\,F_{curr}$.

An example of frequency change detection is presented in Figure 1, where it is possible to observe F_{up}, $\overline{F}_{\max}$ and F_{down}. A change is detected when the maximum frequency of the signal F_{curr} overcomes one of the two thresholds (the horizontal dotted lines):

$$th_{up} = \left(\overline{F}_{\max} + F_{up}\right)/2$$

$$th_{down} = \left(\overline{F}_{\max} - F_{down}\right)/2$$

for h consecutive samples. The choice of h is critical to the robustness of the algorithm: with low values of h (e.g. 1 or 2), the algorithm quickly detects a variation in the maximum frequency of the signal but it might suffer from false detections which can cause a continuous change of the sampling frequency. On the contrary, very high values of h (e.g. 1000 or 2000) decrease the false alarm rate but the algorithm might be less prompt in detecting the changes. We suggest to fix h *=40* in this application to trade off robustness and quickness in the frequency change detection. If available, a-priori information about the process could provide the designer with a suitable parameter h.

The proposed algorithm is summarized in Algorithm 1, where *L* represents the total number of samples in the dataset (we suggest c *=2.1*, h *=40* and W *=400*).

Algorithm 1: Adaptive Sampling Algorithm

1. Estimate $\overline{F}_{\max}$ by considering W initial samples and set $F_c = c\,\overline{F}_{\max}$;
2. Define $F_{up} = \left(1 + \frac{c-2}{4}\right)\overline{F}_{\max}$
   $F_{down} = \left(1 - \frac{c-2}{4}\right)\overline{F}_{\max}$
3. h_1 *=0* and h_2 *=0;*
4. **for** (i=W +1; i < *L*; i++) {
5. Estimate the current maximum frequency F_{curr} on sequence (i-W , i)
6. **if** $\left(\left|F_{curr} - F_{up}\right| < \left|F_{curr} - \overline{F}_{\max}\right|\right)$ {
7. { $h_1 = h_1 + 1;\ h_2 = 0;$ }
8. **else if** $\left(\left|F_{curr} - F_{up}\right| < \left|F_{curr} - \overline{F}_{\max}\right|\right)$
   { $h_2 = h_2 + 1;\ h_1 = 0;$ }
9. **else** { $h_1 = 0;\ h_2 = 0;$ }
10. **if** ($h_1 > h$) || ($h_2 > h$) {
11. $F_c = c\,F_{curr}$;
12. $F_{up} = \left(1 + \frac{c-2}{4}\right)F_{curr}$;
13. $F_{down} = \left(1 - \frac{c-2}{4}\right)F_{curr}$; } }

4. Simulation Setup

We carried out a preliminary simulation analysis to assess the performance of the Adaptive Sampling Algorithm. A simple network scenario consisting of a cluster of sensor nodes equipped with snow sensor boards was considered; a star topology was envisaged for communication towards the base station. Due to its computation demand, the Adaptive Sampling Algorithm is executed at the BS level. The latter computes the optimal sample rate for each single sensor node, based on the data received, and communicates it back to the node. Communication

between sensor nodes and BS occurs in TDMA; this allows for energy saving as the radio is switched off during slots allocated to other nodes.

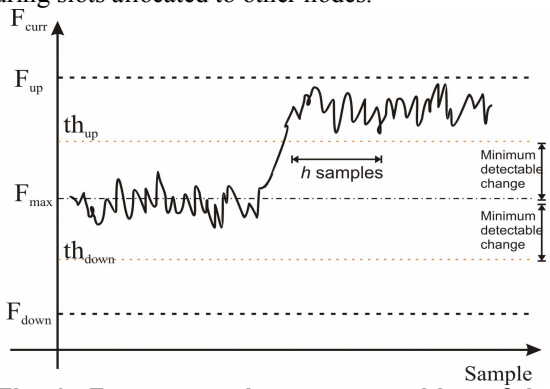

Fig. 1. Frequency change recognition of the adaptive sampling algorithm

We tested our algorithm by using four different data sets, derived from real snow measurements in different days and conditions. Each data set consists of approximately 6.000 samples acquired with a fixed period of 15s. This value was chosen on the basis of a-priori knowledge of the signals to be measured (snow capacitance and ambient temperature): it is large enough to capture quick variations (we should note that it is larger than necessary since we expect snow capacitance and ambient temperature to have small variations for most of the time).

To evaluate the performance of the algorithm we considered the following performance metrics.

- *Sampling Fraction*, defined as the number of samples acquired by the Adaptive Sampling Algorithm divided by the number of samples acquired using fixed-rate over-sampling (i.e., sampling every 15 s).
- *Mean Relative Error (MRE)*, defined as

$$MRE = \frac{1}{N}\sum_{i=1}^{N}\frac{\left|\overline{x_i} - x_i\right|}{\left|x_i\right|}$$

where x_i denotes the $i-th$ sample in the original data sequence, and $\overline{x_i}$ the $i-th$ data sample in the data sequence reconstructed at the BS, and N the total number of data in the original data sequence.

The Sampling Fraction provides an indication of the energy saved by the Adaptive Sampling Algorithm, while MRE gives a measure of the relative error introduced in the data sequence reconstructed at the BS. When the sampling rate estimated by the algorithm is larger than the fixed over-sampling rate some samples in the original sequence are skipped and, thus, not transmitted to the BS. In addition, when the message loss rate is greater than zero, some samples transmitted by the sensor node are missed by the BS. In our experiments we assumed that lost samples are replaced by the previous ones (lost compensation). Hence, in the computation of MRE we assumed $\overline{x_i} = x_i$ if the $i-th$ sample is (correctly) received by BS, and $\overline{x_i} = \overline{x_{i-1}}$ otherwise.

In our experiments messages loss was generated according to a Bernoulli distribution; to increase the accuracy of the simulation results we used the replication method with 90% confidence level [18].

5. Simulation Results

We performed a preliminary set of experiments to tune the algorithm parameters W, c, and h (see Section 3). To this end we used four different data sets measured in different acquisition days and conditions (hereafter referred to as Scenario 1 through Scenario 4), and varied the message loss rate from 0% to 90%. We first varied W in the range [100, 1000] samples, in steps of 100. We found that for values lower than 400 the MRE is very large and exhibits an unstable behavior, while increasing W beyond 400 and up to 1000 increases the energy consumption (more samples are acquired) but does not provide a significant MRE reduction. Then we varied the parameter c in the range [2.1, 3.0]. We observed that a value greater than 2.1 increases the sampling rate but it does not reduce significantly the MRE. Finally, we investigated the behavior of h. We varied h between 10 and 100 and observed the best tradeoff between number of samples generated and MRE when h=40. Based on the above results, in all the subsequent experiments we will use as parameter settings W=400, c=2.1, h=40.

Figure 2 shows the percentage of samples acquired by the snow sensor with adaptive sampling - with respect to the case of fixed over-sampling (i.e., 1 sample every 15 s) - for increasing values of the message loss rate. The Adaptive Sampling Algorithm reduces significantly the number of samples the snow sensor has to acquire. The percentage varies from 18 to 27% depending on the scenario taken into consideration. As a result, the algorithm can save approximately 73-82% of the energy consumed when operating at fixed over-sampling. Moreover, this percentage does not depend significantly on the message loss rate of the wireless link between the sensor node and the base station.

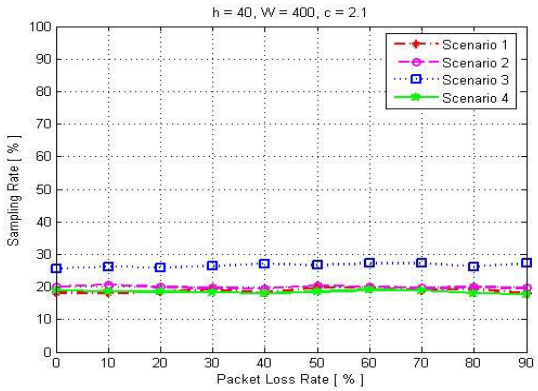

Fig. 2. Sampling Fraction as a function of message loss.

Figure 3 and Figure 4 show the MRE for snow capacitance at low and high frequencies, respectively, while Figure 5 reports the same index for the temperature. The MRE for LF capacitance is under 3% even when the message loss is extremely high. As far HF capacitance, MRE remains always under 4% for all scenarios except in Scenario 4. This is because the corresponding data sequence exhibits several spikes that are filtered out by the algorithm. Nevertheless, the data sequence reconstructed at the BS is always very close to the original data sequence (see Figure 6).

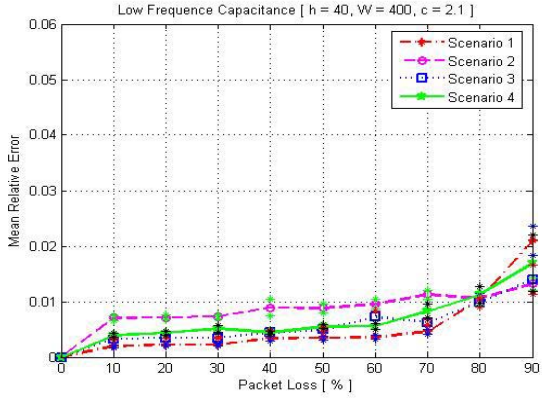

Fig. 3. MRE for low frequency capacitance as a function of message loss.

The MRE for ambient temperature is high in all the scenarios (see Figure 5). This is because temperature values ranges from -3 to 23 C (measurements have been done during spring time, and each data set comes from about 24-hours observation), but it remains close to zero for a large fraction of time (see Figure 7).

When the absolute value is small, small deviations cause high error too. However, Figure 7 shows that the temperature data sequence, which is reconstructed at the BS, is very close to the original one even when the message loss is very high.

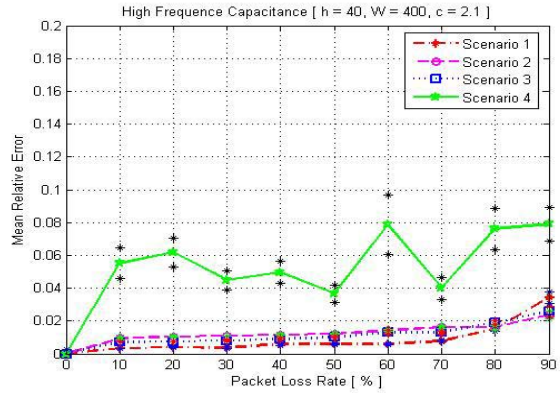

Fig. 4. MRE for high frequency capacitance as a function of message loss.

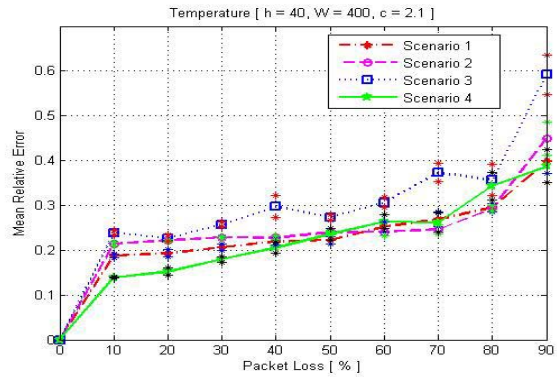

Fig 5. MRE for temperature as a function of message loss.

Summarizing, the Adaptive Sampling Algorithm is able to reduce the number of samples by 73-82% with respect to fixed over-sampling. When used in combination with a simple duty cycle technique that switches off the sensor between consecutive readings (as briefly described in Section 2), the result is a savings of the 95-97% of the energy consumed when the sensor is always kept on.

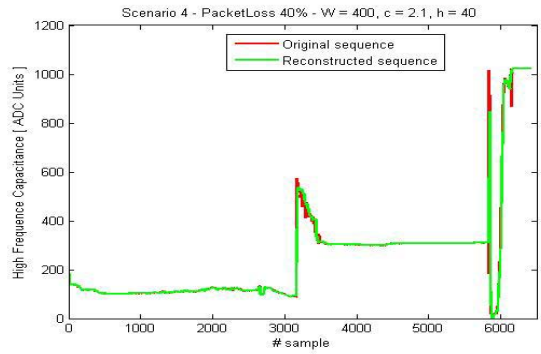

Fig. 6. Original and reconstructed high-frequence capacitance in Scenario 4. Message loss = 40%.

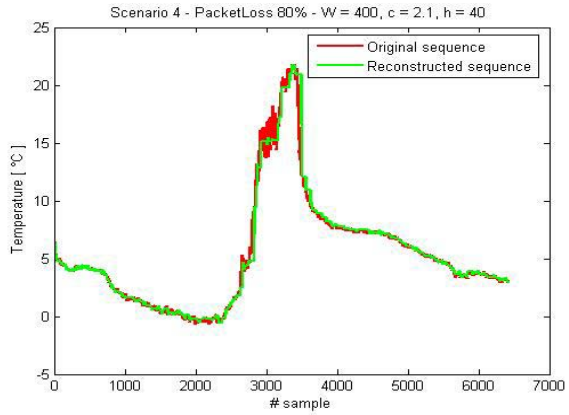

Fig. 7. Original and reconstructed temperature in Scenario 4. Message loss = 80%.

Moreover, the MRE remains at acceptable values even when the message loss is extremely high. In addition, it must be emphasized that, by decreasing the number of samples, the Adaptive Sampling Algorithm reduces accordingly the amount of data to be transmitted by the sensor node, hence reducing the energy for communication by the same percentage (73-82%).

6. Conclusions and Future Work

In this paper we propose an Adaptive Sampling Algorithm to dynamically estimate the optimal sampling frequency of a physical quantity to be monitored over time. The algorithm has been conceived to reduce the energy consumption of a prototype sensor for snow monitoring applications. However, it can be used in all cases where the process to be monitored exhibits slow variation over time. Our simulation experiments have shown that, when used in combination with a duty cycle technique, the above algorithm can save up to 97% of the energy consumed by the snow sensor when it is always on. In addition, by reducing the amount of data to be transmitted, the algorithm also reduces significantly the energy consumed by the radio. We have analyzed our proposal in a very simple network scenario. The next step will consist in exploring more complex network scenarios.

7. References

[1] I.F.Akyildiz, W. Su, Y. Sankarasubramaniam E. Capirci, “Wireless Sensor Networks: a Survey”, Computer Networks, Vol 38, N. 4, March 2002.

[2] C. Tang, C. Raghavendra, “Compression Techniques for Wireless Sensor Networks”, Chapter 10 in Wireless Sensor Networks (C. Raghavendra, K. Sivalingam, T. Znati, Eds.), Kluwer Academic Publishers, Boston, USA, 2004.

[3] C. Sadler, M. Martonosi, “Data Compression Algorithms for Enegy-Constrained Devices in Delay Tolerant Networks”, Proc. ACM SenSys 2006, Boulder, Colorado, USA, November 2006..

[4] S. Madden, M. Franklin, J.Hellerstein, W. Hong. “TAG: a Tiny AGgregation Service for Ad-Hoc Sensor Networks”, Proceedings of OSDI, 2002.

[5] A. Boulis, S.Ganeriwal, M. B. Srivastava, “Aggregation in Sensor Networks : an Energy-Accuracy Trade-off ”, Ad Hoc Networks, Vol. 1 (2003), pp. 317-331.

[6] S. Goel, T. Imielinsky, “Prediction-Based Monitoring in Sensor Networks: Taking Lessons from MPEG”, ACM Computer Communication Review 2001, Vol. 31, N. 5, October 2001.

[7] A. Cerpa, D. Estrin, “ASCENT: Adaptive Self-Configuring Sensor Networks Topologies”, Proc. IEEE Infocom 2002, New York, USA, 2002.

[8] C. Schurgers, V. Tsiatsis, S. Ganeriwal, M. Srivastava, “Optimizing Sensor Networks in the Energy-Latency-Density Design Space”, IEEE Trans. on Mobile Computing, Vol.1, N.1, 2002.

[9] D. Ganesan, A. Cerpa, W. Ye, Y. Yu, J. Zhao, D. Estrin,, “Networking Issues in Wireless Sensor Networks”, Journal of Parallel and Distributed Computing, Vol. 64 (2004).

[10] V. Raghunathan, S. Ganeriwal, M. Srivastava, "Emerging Techniques for Long Lived Wireless Sensor Networks", IEEE Communications Magazine, April 2006, pp. 108-114.

[11] G. Anastasi, M. Conti, M. Di Francesco, A. Passarella, “How to prolong the Lifetime of Wireless Sensor Networks” in Mobile Ad Hoc and Pervasive Communications, (M. Denko, L. Yang, Eds.), to appear.

[12] A. Jain, E. Y. Chang, “Adaptive Sampling for Sensor Networks”, Proc. Workshop on Data Management for Sensor Networks (DMSN 2004), Toronto (CA), 2004.

[13] R. Willet, A. Martin, R. Nowak, “Backcasting: Adaptive Sampling for Sensor Networks”, Proc. IPSN 2004, April 26-27, 2004, Berkeley, USA.

[14] J. Zhou, D. De Roure, S. Vivekanandan, “Adaptive Sampling and Routing in a Floodplain Monitoring Sensor Network”, Proc. IEEE WiMob 2006, June 19-21, 2006

[15] M. Basseville, and I.V. Nikiforov, “Detection of Abrupt Changes: Theory and Application”, Prentice-Hall, Inc. 1993.

[16] M. Niang, M. Bernier, Y. Gauthier, G. Fortin, E. Van Bochove, M. Stacheder, and A. Brandelik, “On the validation of snow densities derived from snowpower probes in a temperate snow cover in eastern canada: First result”, in 60th Eastern Snow Conference, Québec (CA), 2003, pp. 176-187.

[17] A. J. Jerri, "The Shannon Sampling Theorem—Its Various Extensions and Applications: A Tutorial Review", Proc. of IEEE, 65:1565–1595, Nov. 1977.

[18] J. Banks, “Handbook of Simulation”, John Wiley and Sons, New York, USA, 1998.