An Adaptive Sampling Algorithm for Effective Energy Management in Wireless Sensor Networks with Energy-hungry Sensors

Cesare Alippi[*], Giuseppe Anastasi[§], Mario Di Francesco[§], 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 wireless sensor networks generally assume that data acquisition and processing have an energy consumption significantly lower than that of communication. Unfortunately, this assumption does not hold in a number of practical applications, where sensors may consume even more energy than the radio. In this context, effective energy management should include policies for an efficient utilization of the sensors, which become one of the main components affecting the network lifetime. In this paper we propose an *Adaptive Sampling Algorithm* that estimates on line the optimal sampling frequencies for sensors. This approach, which requires the design of adaptive measurement systems, minimizes the energy consumption of the sensors and, incidentally, also that of the radio, while maintaining a very high accuracy of collected data. As a case study, we considered a sensor for snow monitoring applications. Simulation experiments have shown that the suggested adaptive algorithm can reduce the number of acquired samples up to 79% with respect to a traditional fixed-rate approach. We have also found that it can perform similar to a fixed-rate scheme where the sampling frequency is known in advance.

1. Introduction

Wireless sensor networks (WSNs) are distributed measurement systems consisting of a large number of measurement units deployed over a geographical area; each unit is a low-power device that integrates processing, sensing and wireless communication abilities. Units acquire information from the surrounding environment and, after (a possible) local processing, send measurements to one or more collection points or base stations for further data aggregation and interpretation [1].

Among the set of potential scenarios, monitoring applications can particularly benefit from this technology as WSNs allow a long-term data collection at scales and resolutions that are difficult, if not impossible, to achieve with traditional techniques [2]. In recent years, the number of WSN deployments for real life applications has rapidly increased and this trend is expected to increase even more in next years [3], [4]. However, energy consumption still remains the major

obstacle for the full diffusion and exploitation of this technology, even when batteries can be recharged, e.g., through solar energy harvesting mechanisms [5].

Table 1. Power consumption for some common radios [6].

Radio	Producer	Power Consumption	
		Transmission	Reception
CC2420	Texas Instruments	35 mW (at 0 dBm)	38 mW
CC1000	Texas Instruments	42 mW (at 0 dBm)	29 mW
TR1000	RF Monolithics	36 mW (at 0 dBm)	9 mW

Table 2. Power consumption for some off-the-shelf sensors.

Sensor	Producer	Sensing	Power Consumption
STCN75	STM	Temperature	0.4 mW
QST108KT6	STM	Touch	7 mW
SG-LINK (1000Ω)	MicroStrain	Strain gauge	9 mW
iMEMS	ADI	Accelerometer (3 axis)	30 mW
2200 Series, 2600 Series	GEMS	Pressure	50 mW
T150	GEFRAN	Humidity	90 mW
LUC-M10	PEPPERL+FUCHS	Level Sensor	300 mW
TDA0161	STM	Proximity	420 mW
FCS-GL1/2A4-AP8X-H1141	TURCK	Flow Control	1250 mW

In last years, many energy conservation schemes have been proposed in the literature (a detailed survey can be found in [7]), which assume that data acquisition and processing have an energy consumption significantly lower than communication (as a consequence, the research aims at minimizing the radio activity). Only recently, the progressive utilization of distributed measurement systems for monitoring complex phenomena has shown that the above assumption does not necessarily hold. In fact, many real life applications require specific sensors whose power consumption cannot be neglected [8]. Table 1 and Table 2 provide the power consumptions of the most popular radio equipments used in sensor nodes and some off-the-shelf sensors, respectively (the selection of sensors has been made only to ease the reader's understanding). If we also consider that acquisition times are typically longer than transmission ones, we can conclude that some sensors may even consume significantly more energy than the radio.

As such, energy conservation schemes aiming at minimizing the radio activity need to be complemented with techniques implementing an efficient energy management of the sensors.

In this paper we propose a general approach that leverages two complementary mechanisms at the sensor level: (**i**) *duty cycling* (i.e., the sensor board is switched off between two consecutive samples); and (**ii**) *adaptive sampling* (i.e., the optimal sampling frequency is estimated online).

In particular, we suggest an *Adaptive Sampling Algorithm* (ASA) that adapts the sampling frequencies of the sensors to the evolving dynamics of the process.

In the Instrumentation and Measurement community the adaptive sampling approach has been applied to address several issues. For instance, [9] has suggested an adaptive sampling technique for measuring the difference in phase between the fundamental components of two signals: the sampling rate is increased until the phase is correctly measured or the sampling rate reaches the maximum sampling rate of the system.

[10] describes a Fourier analyzer which autonomously adapts the parameters of the filters to match the signal components and the measuring channels. The result is that the picket-fence effect and leakages are reduced (but the method can be applied only to periodic signals). [11] presents a velocity adaptive measurement system for closed-loop position control which relies on the adaptation of the sampling frequency to improve the response time.

In [12] the authors propose a decentralized approach to adaptive sampling which uses a Kalman filter to predict the sensor node activity and adjust the sampling frequency correspondingly.

ASA is more general than the above solutions since it does not assume any hypothesis regarding the nature of the signal (e.g., stationarity); moreover, its computational load is acceptable for mid complexity WSN units.

Since ASA identifies on-line the minimal sampling frequency guaranteeing reconstruction of the sampled signal it reduces the power consumption of the measurement phase by adapting the sampling frequency to the real needs of the physical phenomena under observation.

By decreasing the number of acquired samples ASA also reduces the amount of data to be transmitted and, as a consequence, the energy consumed by the radio. In addition, the proposed approach can be integrated with other techniques for energy conservation acting at different abstraction levels (e.g., data aggregation and/or compression).

The paper is organized as follows. Section 2 introduces the Adaptive Sampling Algorithm. Section 3 presents the snow sensor that is used as a case study to assess the performance of the algorithm and the simulation environment used for performance analysis. Simulation results are finally presented in Section 4.

2. Adaptive Sampling Algorithm: the proposed methodology

The proposed ASA algorithm differentiates over the above presented literature by adapting online the sampling frequency of the sensor to the physical phenomenon under monitoring provided that a change in the maximum frequency is detected.

Detecting a change in a noisy nonstationary environment is quite an open research issue generally addressed with a statistical approach e.g., see [13][14][15]. Here, we found particularly appropriate a modification of the CUmulative SUM (CUSUM) change detection test [16], widely used in the system control community. In particular, we configured the test to detect changes associated with the highest frequency F_{max} of the signal, being F_{max} related to the minimum sampling frequency F_N as per Nyquist $F_N > 2\ F_{max.}$ [17].

Frequency F_{max} is not available a priori and changes over time in a nonstationary process. Consequently, it clearly emerges that F_N changes over time as well and that, by adapting the sampling frequency, over-sampling is avoided, signal reconstruction guaranteed and power consumption reduction obtained.

The proposed algorithm initially estimates, through a Fast Fourier transform, $\overline{F}_{\max}$ by using the first W acquired data which are assumed to be generated by a stationary process. The initial sampling frequency is $F_c = c\,\overline{F}_{\max}$, where c is a confidence parameter that, according to Nyquist, must be larger than 2 (it is common to pick a sampling frequency three to five times higher than the signal maximum frequency [18]). The value of $\overline{F}_{\max}$ can be determined with different techniques. Here, $\overline{F}_{\max}$ is identified by relying on a Signal to Noise ratio philosophy: $\overline{F}_{\max}$ is the frequency for which the ratio between the energy of the signal up to $\overline{F}_{\max}$ and the energy of the residual segment of the right spectrum starting from $\overline{F}_{\max}$ is 100.

To allow CUSUM for detecting the change in the maximum frequency we designed the two alternative hypotheses

$$F_{up} = \min\left\{(1+\delta)\cdot \overline{F}_{\max}, \frac{F_c}{2}\right\}; \quad F_{down} = (1-\delta)\cdot \overline{F}_{\max},$$

which address an increment and a decrement in the maximum frequency, respectively. $\delta \in \Re^+$ is an user-defined confidence parameter which represents the *minimum detectable frequency change*. In other words, δ represents the minimum percentage change in the maximum frequency which must be detected by ASA. Of course, this is an independent parameter which must be set by the user/designer, e.g., δ=3% implies that changes affecting $\overline{F}_{\max}$ for more than 3% $\overline{F}_{\max}$ must be detected (to be intended in statistical terms). F_{up} might be significantly influenced by the value of δ. In fact, when δ is small then $F_{up} \cong \overline{F}_{\max}$ and we should expect an increment in the number of false positives in detection. Vice versa, for higher δs, $F_{up} \cong \frac{F_c}{2}$ and the algorithm might suffer from the presence of false negatives. Obviously, F_{up} cannot be larger than $\frac{F_c}{2}$ due to the Nyquist theorem. F_{down} is less influenced by δ since a decrement in the maximum frequency above $\delta\overline{F}_{\max}$ is detected (aliasing effects are here not introduced).

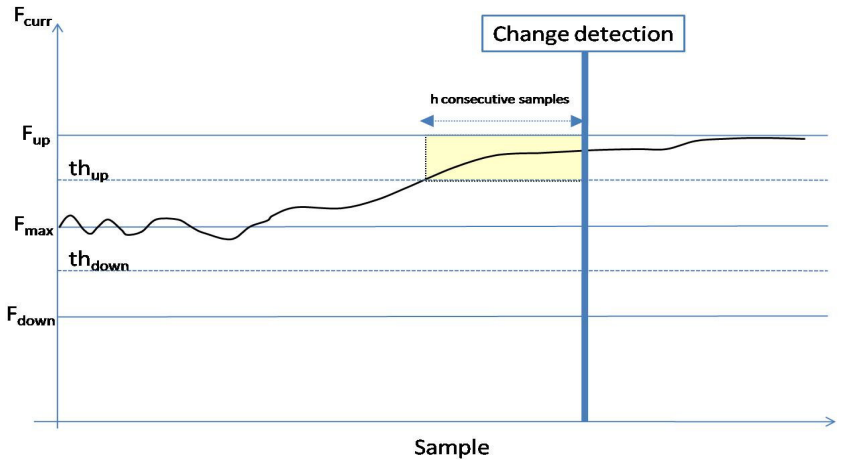

Figure 1. Detecting a change in the maximum frequency.

During operational life, a change is detected in the process when the current maximum frequency F_{curr} (estimated over a W samples sequence) overcomes one of the thresholds:

$$th_{up} = \frac{\overline{F}_{\max} + F_{up}}{2} = \overline{F}_{\max}(1+\delta/2),$$

$$th_{down} = \frac{\overline{F}_{\max} + F_{down}}{2} = \overline{F}_{\max}(1-\delta/2)$$

for h consecutive samples. An example of a frequency change is illustrated in **Figure 1**. When a change in the maximum frequency is detected the sampling frequency is modified according to the new value to track the process evolution. In short, ASA can be synthesized in the detection rule: *if* $\left(\left|F_{curr} - F_{up}\right| < \left|F_{curr} - \overline{F}_{\max}\right|\right)$ *for* h *consecutive samples or if* $\left(\left|F_{curr} - F_{down}\right| < \left|F_{curr} - \overline{F}_{\max}\right|\right)$ *for* h *consecutive samples, then the new sampling frequency is* $F_c = c\,F_{curr}$.

The proposed algorithm is given in Algorithm 1.

Algorithm 1: Adaptive Sampling Algorithm (c, δ, h)

1. Store the W initial samples coming from the process in *Dataset;*
2. Estimate $\overline{F}_{\max}$ on *Dataset* and set $F_c = c\,\overline{F}_{\max}$;
3. Define $F_{up} = \min\left\{(1+\delta)\cdot\overline{F}_{\max}, \frac{F_c}{2}\right\}$; $F_{down} = (1-\delta)\cdot\overline{F}_{\max}$;
4. $h_1 = 0$ and $h_2 = 0$; $i = W + 1$;
5. **while** (1) {
6. Acquire the *i*-th sample and add it to *Dataset;*
7. Estimate the current maximum frequency F_{curr} on the sequence *Dataset*(i-W +1, i);
8. **if** $\left(\left|F_{curr} - F_{up}\right| < \left|F_{curr} - \overline{F}_{\max}\right|\right)$
9. $h_1 = h_1 + 1$; $h_2 = 0$;
10. **else if** $\left(\left|F_{curr} - F_{down}\right| < \left|F_{curr} - \overline{F}_{\max}\right|\right)$
$h_2 = h_2 + 1$; $h_1 = 0$;
11. **else** $h_1 = 0$; $h_2 = 0$;
12. **if** ($h_1 > h$) || ($h_2 > h$) {
13. $F_c = c\,F_{curr}$;
14. $F_{up} = \min\left\{(1+\delta)\cdot\overline{F}_{\max}, \frac{F_c}{2}\right\}$;
15. $F_{down} = (1-\delta)\cdot\overline{F}_{\max}$;
16. $\overline{F}_{\max} = F_{curr}$;
17. }
18. }

Low values of h (e.g., 1 or 2) allow the algorithm for quickly detecting a variation in the maximum frequency of the signal (but we could experience false positives inducing a continuous

change of the sampling frequency). On the contrary, high values of h (e.g., 1000 or 2000) decrease the false alarm rate at the expenses of a slower promptness in detecting the change. The value of h can be either user-defined (e.g., by exploiting available a priori information about the process) or estimated by ASA as suggested in Algorithm 2.

Define T as the number of initial stationary samples used to configure h (T must be sufficiently large to grant that the estimate of h converges towards its expected value): we suggest as estimate of h the count of the maximum number of subsequent false positives in the training sequence.

Algorithm 2: h = **Automatic Configuration of ASA** (c, δ, W, T)

---

1. Estimate $\overline{F}_{\max}$ by considering W initial samples and set $F_c = c\,\overline{F}_{\max}$;
2. Define $F_{up} = \min\left\{(1+\delta)\cdot\overline{F}_{\max},\ \frac{F_c}{2}\right\}$; $F_{down} = (1-\delta)\cdot\overline{F}_{\max}$;
3. $h_1 = 0$ and $h_2 = 0$;
4. $\tilde{h}_1 = 0$ and $\tilde{h}_2 = 0$;
5. **for** (i=W+1; i < T; i++) {
6. Estimate the current maximum frequency F_{curr} on sequence (i-W+1, i)
7. **if** $\left(\left|F_{curr} - F_{up}\right| < \left|F_{curr} - \overline{F}_{\max}\right|\right)$
8. $h_1 = h_1 + 1;\ h_2 = 0;$
9. **else if** $\left(\left|F_{curr} - F_{down}\right| < \left|F_{curr} - \overline{F}_{\max}\right|\right)$
$h_2 = h_2 + 1;\ h_1 = 0;$
10. **else** $h_1 = 0;\ h_2 = 0;$
11. **if** ($h_1 > \tilde{h}_1$) {
12. $\tilde{h}_1 = h_1;$
13. }
14. **if** ($h_2 > \tilde{h}_2$) {
15. $\tilde{h}_2 = h_2;$
16. }
17. }
18. **return** $\left(\min\left(\tilde{h}_1, \tilde{h}_2\right)\right)$;

More in detail, Algorithm 2 operates as follows. At first $\overline{F}_{\max}$ is estimated on the initial W samples of the training sequence (line 1), then F_{up} and F_{down} are computed according to ASA (line 3 of Algorithm 1). The procedure (line 7 - 16) counts the maximum number of consecutive samples in subsequent samples for which F_{curr} is closer either to F_{up} (the counter is $\widetilde{h}_1$) or F_{down} (the counter is $\widetilde{h}_2$). In order to be conservative h is the minimum between $\widetilde{h}_1$ and $\widetilde{h}_2$ (line 18).

The ASA algorithm runs at the base station which notifies updates of the current sampling frequency to remote units (the algorithm might be too complex to be executed on tiny devices). However, from the conceptual point of view, there would be no objection in using a decentralized approach which executes ASA at the sensor node level.

3. Experimental setup

3.1 Description of the snow sensor

To evaluate the performance of ASA in a real application we considered a sensor developed for monitoring the snow composition (slope stability assessment and avalanches forecast). Such a sensor provides the dataset used in subsequent simulation experiments.

The snow sensor considered in this work is a multi-frequency capacitive measuring unit engineered to be embedded in a remote wireless measuring system. It is composed of a probe, a main multi-frequency injection board capable of measuring capacity at different frequencies [19], and a wireless unit to be left on the mountain (for example fixed on a pole); the system is powered by a rechargeable battery pack.

At each sampling cycle the snow sensor provides measurements of snow capacitance at 100 Hz (low frequency) and 100 kHz (high frequency); such frequencies have been proved to differentiate water from air, snow and ice. At the same time a second sensor provides a measurement of the ambient temperature.

The snow capacitances at low and high frequency and the temperature information is passed to the sensor node, packed in a single message and sent over the wireless channel. For each measurement the electronic injection board 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; such a choice leads to approximately 150 mJ energy consumption per sample. When the duty cycle mechanism substitutes the fixed sampling approach an immediate energy saving arises (here, the energy consumption of the sensor decreases of about 80%.)

3.2 Sensor network configuration

ASA can be implemented in any sensor network architecture. However, to the purpose of simulation, we considered a cluster-based architecture (see Figure 2). For each node, the sampling frequency is computed and dynamically updated at the base station and, then, notified to the node through special notification messages. We also defined a communication protocol similar to LEACH [20] for collecting data from nodes to the base station and diffusing sampling frequency notifications in the back direction (details are given in [21]). We implemented both ASA and the cluster-based communication protocol by using the TOSSIM simulation tool [22], a widely used simulator for WSNs. In the considered communication protocol, both nodes within a cluster and cluster heads use a TDMA scheme for exchanging data with the corresponding cluster head or base station, respectively. Each node (cluster head) remains active only during the time slots assigned for communication so as to minimize the radio energy consumption. Intra-cluster interferences (i.e., collisions due to simultaneous transmissions of nodes belonging to the same cluster) are thus avoided by the communication protocol, while inter-cluster interferences (i.e., collisions due to simultaneous transmissions by nodes belonging to different clusters) can still occur. We modeled the effects of possible inter-cluster interferences as message losses. Therefore, in our simulations, messages may be missed either due to transmission errors or inter-cluster interferences.

The communication protocol uses an ARQ (Automatic Repeat Request) scheme based on acknowledgments, timeouts and retransmissions to recover missed messages. Messages not acknowledged within the timeout time are retransmitted up to a predefined maximum number of times (see [21] for details). In case of missing samples (i.e., messages which did not reach the base station after the retransmission), the base station uses a simple loss compensation technique

by replacing a missing sample with the previous one. This is a very simple approach which, nevertheless, proves to be effective in increasing the accuracy of the data sequence collected at the base station even when the wireless communication is not completely reliable. Of course, alternative, more complex, loss compensation schemes, e.g., based on data missing reconstruction, can be considered within the proposed framework.

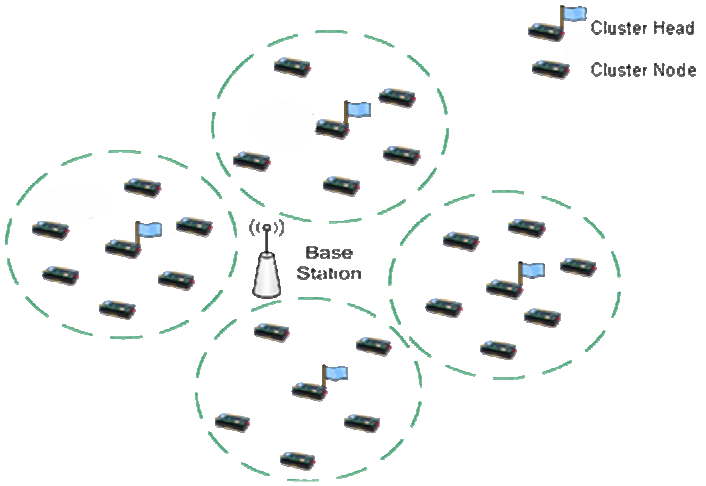

Figure 2. Cluster-based sensor network architecture.

3.3 Figures of merit

To measure the performance of ASA we defined the following figures of merit.

- *Sampling Fraction*, defined as the number of samples acquired by the sensor according to ASA w.r.t. the number of samples acquired with a fixed sampling frequency. The Sampling Fraction aims at evaluating the efficiency of the algorithm.
- *Sensor/Radio Energy Consumption,* which summarizes the total energy consumed by the sensor/radio subsystem. The total energy consumption of the sensor is the product between the energy drained by each sampling cycle and the number of samples generated during the simulation experiment. The total energy consumed by the radio can be modeled as

$$E_R = T_t \cdot P_t + T_r \cdot P_r + T_i \cdot P_i + T_s \cdot P_s .$$

T_t, T_r, T_i, T_s represent the total time the radio is in the transmitting, receiving, idle, or sleeping operational modes while P_t, P_r, P_i, P_s refer to the associated power consumptions. We assume in our simulations that $P_t \approx P_r \approx P_i$, and that the power consumed in the sleeping mode is negligible w.r.t. the power consumed in the others modalities. Therefore, in our

simulator, we implemented the approximated model

$$E_R = T_a \cdot P_a$$

where T_a denotes the total time during which the radio is active (i.e., $T_a = T_t + T_r + T_i$) and $P_a = max(P_t, P_r, P_i)$.

- *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, $\overline{x_i}$ the $i-th$ sample in the data sequence reconstructed at the base station and N the total number of samples in the original data sequence, respectively. MRE gives a measure of the relative error introduced by the algorithm in the data sequence reconstructed at the base station. To measure the accuracy of the temperature sequence we also considered the Mean Absolute Error, which provides better indications than MRE in this specific case. It is defined as

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

3.4 Parameter Settings and Methodology

In our simulations we assumed that nodes are equipped with the Chipcon CC1000 radio (used in the MICA2 motes series) whose operating parameters (derived from [6]) are shown in Table 3. To set the parameters of ASA we referred to a preliminary analysis carried out in a previous paper [23] suggesting W=512, c=2.1, $\delta = 2.5\%$, h=40. Finally, the parameter values used by the communication protocol for collecting data and diffusing sampling-rate notifications are reported in [21].

Table 3. Radio Parameters.

Parameter	Value
Radio	CC1000
Frame size	36 bytes
Bit rate	19.2 Kbps
Transmit Power (0 dBm)	42 mW
Receive Power	29 mW
Idle Power	29 mW
Sleeping Power	0.6 μW

We assessed the performance of ASA by using four different datasets derived from real measurements with the snow sensor described in Section 3.1 in different days and conditions. Each dataset consists of approximately 6.000 samples acquired with a fixed period of 15s. This sampling frequency 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 expected snow capacitance and ambient temperature to change over time).

In the experiments message losses were modeled according to a Bernoulli distribution. To improve the accuracy of the simulation results we used the replication method with 90% confidence level [24].

In the following, when not differently specified, figures refer to experiment 2 being the most critical one for ASA. However, results are similar for other datasets.

4. Simulation Results

We divided our analysis into two parts. At first we investigated the pros and cons of using ASA – in terms of energy saving and impact on the data accuracy– with respect to a fixed sampling-rate approach. Then, we studied the influence of the communication reliability on the performance of the adaptive algorithm.

4.1 Adaptive vs. Fixed Sampling

In the first set of experiments we compared the evolution over time of the current maximum frequency $F_{\max}$ computed over sliding windows of the input signal with that of $\overline{F}_{\max}$ as set by ASA (we remind that $F_c = c\overline{F}_{\max}$).

As presented in Figure 3, we appreciate the fact that ASA is effective in adapting the sampling frequency to the real needs of the physical phenomenon under monitoring. The Figure shows $F_{\max}$ and $F_c/2$ that are the maximum frequency currently available in the signal and the maximum frequency detectable by ASA according to Nyquist's theorem, respectively. Obviously, when $F_c/2 < F_{\max}$ aliasing effects may occur but ASA reacts by increasing the maximum sampling frequency.

Initially, the sampling frequency was set to 1/15Hz up to sample 512 (we did not know the initial optimal sampling rate and we opted for having it overdimensioned). Then, ASA reduces the sampling frequency to 1/75Hz (obviously, any change in F_c is reflected on $F_c/2$); this allows us for reducing the number of acquisitions while maintaining the signal reconstructing ability at the base station (under the no message loss hypothesis). We experience an increase in $F_{\max}$ around sample 620 and ASA adapts F_c to be 1/60Hz. Finally, the increase in F_{curr} around sample 990 induces a next increment in the sampling frequency to 1/45Hz. We note that the abrupt change in frequency occurring at around sample 990 introduces aliasing phenomena since $F_{\max}$ is larger than $\overline{F}_{\max}$. Once nonstationarity has been detected ASA intervenes and adapts the frequency at sample 1030 (with a delay function of the window size W and the change detection mechanism, here the delay is about 40 minutes since the sampling rate is about 1 sample per minute and h=40).

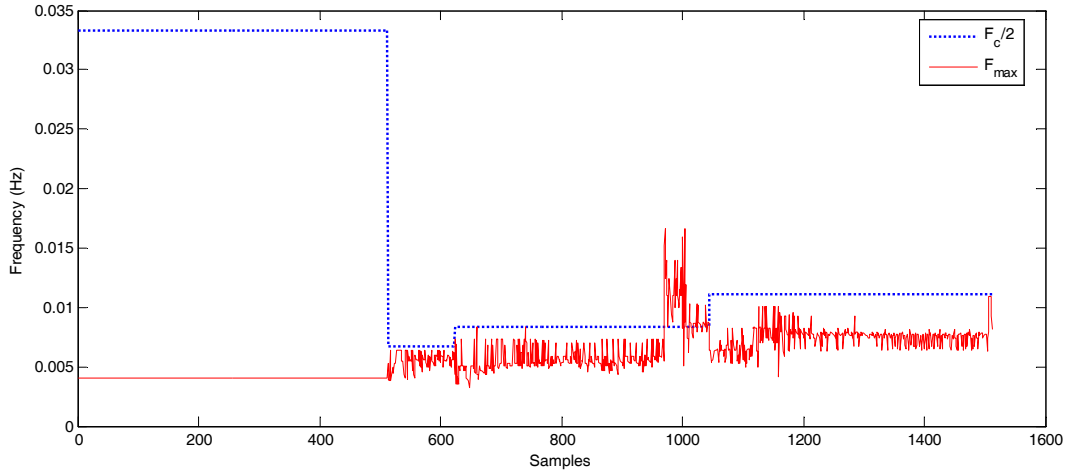

Figure 3 - Sampling rate as a function of the samples.

When messages are lost during communication the general trend does not change apart from a further delay in adapting the sampling frequency (due to the larger delay experienced by notification messages).

We comment that, initially, the sensor node uses the maximum sampling rate 1/15Hz. Afterwards, once the base station has received W=512 samples, the new sampling frequency is

notified to the sensor node. Then, the base station continues to compute the sampling rate based on the received samples so as to adjust it to the current dynamics of the signals.

Figure 4 shows the Sampling Fractions – with respect to both fixed sampling periods (i.e., 15s and the optimal) for the various datasets, and for different values of (hop-by-hop) message loss probability. At each hop, messages missed by the destination are retransmitted up to 2 times. ASA is able to reduce significantly the number of samples with respect to a traditional approach based on a fixed sampling frequency. The number of samples to be acquired is reduced to 21-34% (depending on the dataset and message loss rate) with respect to the 1/15Hz sampling rate. Figure 4 shows that ASA may also outperform the optimal (but unfeasible) fixed-rate approach in terms of Sampling Fraction (and, hence, energy efficiency), especially when the communication is reliable (i.e., for low probability of message loss). This is a consequence of the fact that ASA is able to adapt the sampling frequency to the current signal dynamics, and can thus take advantage of current demands.

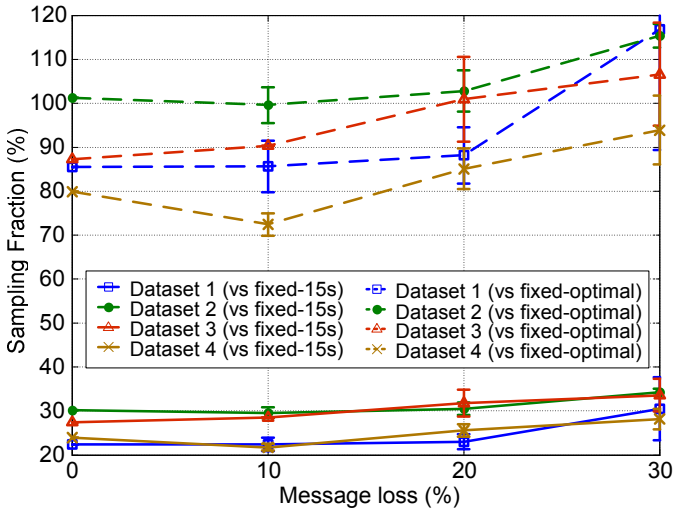

Figure 4. Sampling Fraction as a function of the message loss rate for different datasets.

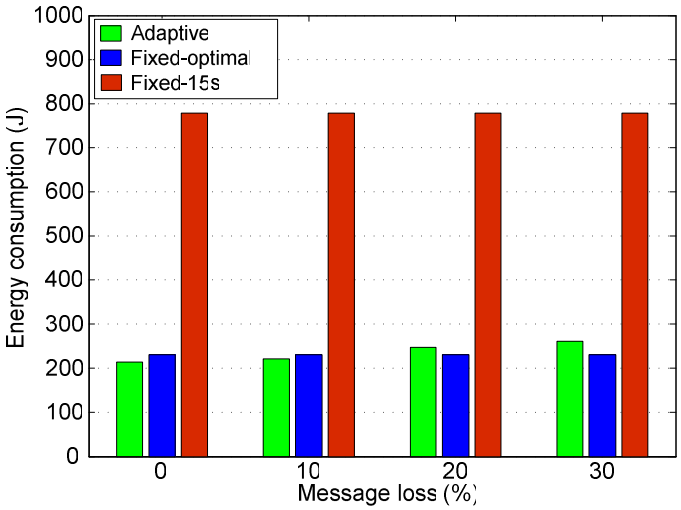

Figure 5. Total energy consumed by the snow sensor.

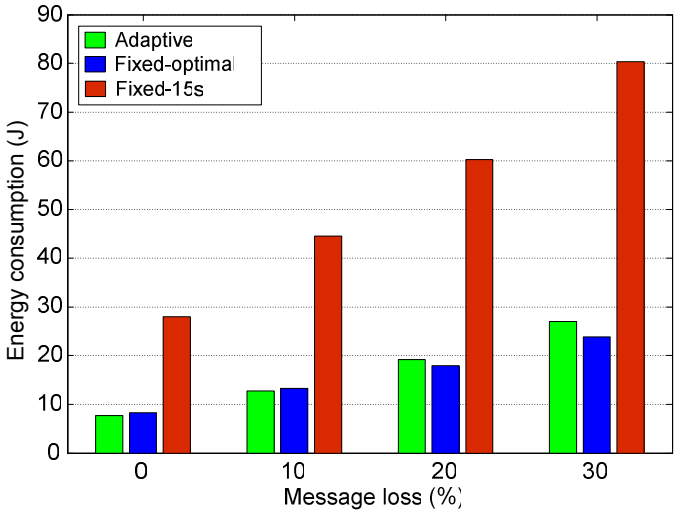

Figure 6. Energy consumed by the radio.

The decrease in the Sampling Fraction provided by ASA results in a corresponding decrease in the energy consumption for both sensing and communication. Figure 5 and Figure 6 show the total energy consumed by the snow sensor and the radio, respectively for dataset 2. Note that the energy savings provided by ASA are not obtained at the expense of a decreased accuracy in the data sequence collected at the base station. Figure 7 and Figure 8 show the MRE for the snow capacitance at low and high frequencies, respectively, while Figure 9 presents the same index for the temperature. The MRE, both at low and high frequencies, remains very low (i.e., 1%-2%) even when the (hop-by-hop) message loss probability increases up to 30%. On the contrary, the MRE for the temperature is high for all datasets (see Figure 9). This is because the temperature

ranges from -3 to 23 °C (measurements have been done during spring time) and remains close to zero for a large fraction of the experiment. Thus, when the absolute value of the measurement is close to zero, even small deviations from such value can cause large errors. Hence, in this specific case we also evaluated the Mean Absolute Error (MAE), summarized in Table 4 for the different datasets and message loss probabilities. We can see that the average (absolute) deviation from the original temperature sequence is always negligible.

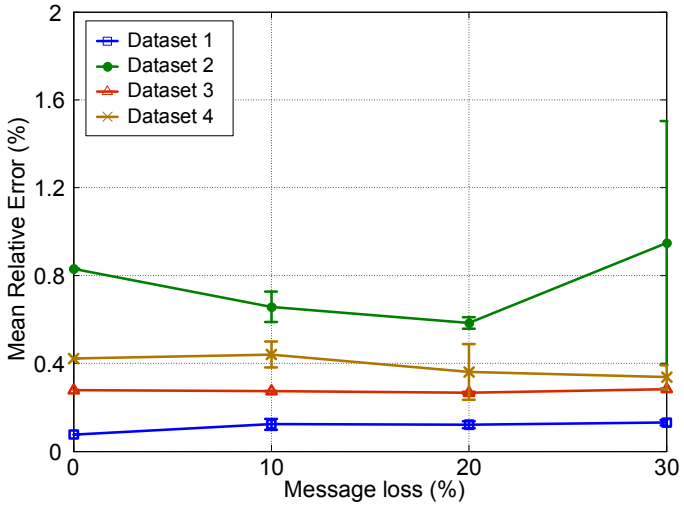

Figure 7. MRE for low frequency capacitance as a function of the message loss.

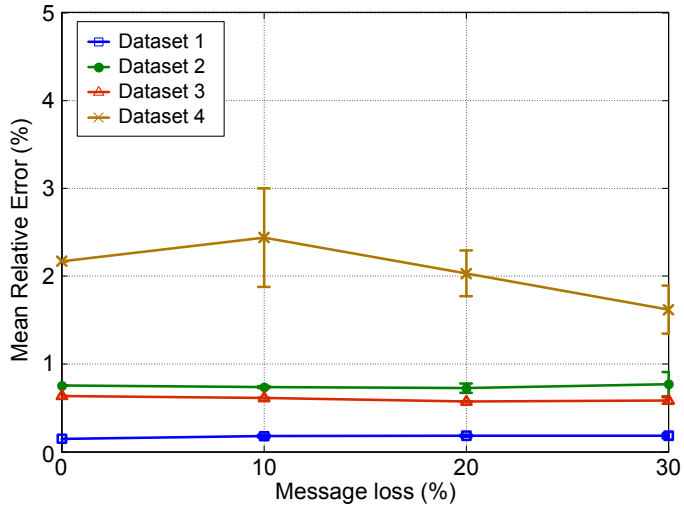

Figure 8. MRE for high frequency capacitance as a function of the message loss.

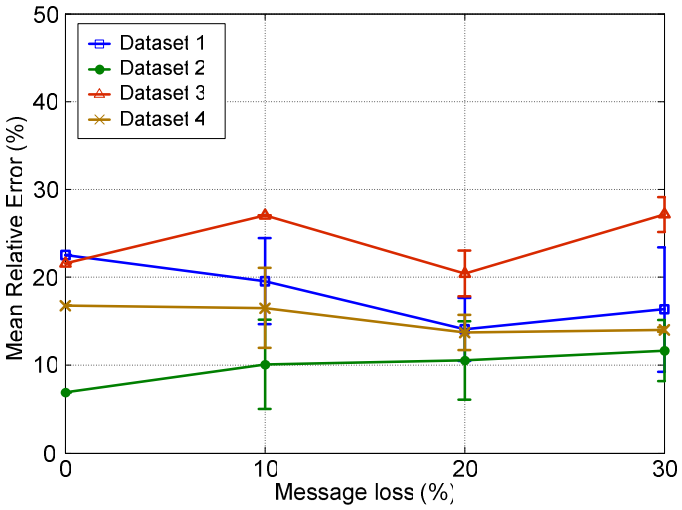

Figure 9. MRE for temperature as a function of the message loss.

	Message loss=0%	Message loss=10%	Message loss=20%	Message loss=30%
Dataset 1	0.06	0.07	0.07	0.07
Dataset 2	0.07	0.08	0.07	0.08
Dataset 3	0.08	0.07	0.06	0.07
Dataset 4	0.09	0.10	0.09	0.10

Table 4. MAE for the temperature as a function of the message loss in °C.

4.2 Impact of communication unreliability

In the previous section emerged that ASA can actually reduce the percentage of samples to be acquired while assuring the information to be delivered at the base station. However, its performance –in terms of energy efficiency– degrades as the (hop-by-hop) message loss probability increases (see Figures 4-6) since ASA reduces the sampling frequency by exploiting the temporal correlation among consecutive samples. As such, to work correctly, ASA requires (almost) all data to be received by the base station. Actually, the algorithm is able to tolerate a certain fraction of missing samples, thanks to the loss compensation mechanism (the phenomenon under monitoring is assumed to change slowly over time). However, when the percentage of missing messages becomes significant, ASA may react by increasing the sampling frequency. Moreover, if the communication is unreliable, notifications sent by the base station to sensor nodes for updating the sampling frequencies may get lost, or experience a large delay. Thus, a node might operate with obsolete sampling frequencies even for a long time. If the

sampling frequency for a sensor is higher than required, oversampling occurs. If lower, aliasing effects may occur.

To make the adaptive sampling approach effective, we should thus guarantee a message delivery ratio (i.e., percentage of messages correctly received by the final destination) in both directions close to 100%, e.g., through an acknowledgment-based retransmission protocol as here considered. In [25] it has been found that retransmission is the most efficient approach to data transfer reliability in wireless sensor networks. Obviously, message retransmission increases the delivery ratio at the cost of additional energy consumption: it is thus important to evaluate the impact of message retransmissions in terms of energy consumption for the overall system (both sensor and radio). A set of experiments was then carried out in which the maximum number of retransmissions per message, max_rtx, changed in the [0,3] range.

Figure 10 shows the impact of the max_rtx value on the Sampling Fraction for increasing message loss probabilities. As expected, the Sampling Fraction decreases significantly when the maximum number of retransmissions per message increases, as the delivery ratio increases accordingly, as shown in Figure 11. With $max\_rtx \geq 2$, more than the 85% of messages are delivered to the final destination (even with a link message loss probability of 30%), and the performance of ASA is similar to, or even better than, that of the (unfeasible) fixed-rate approach.

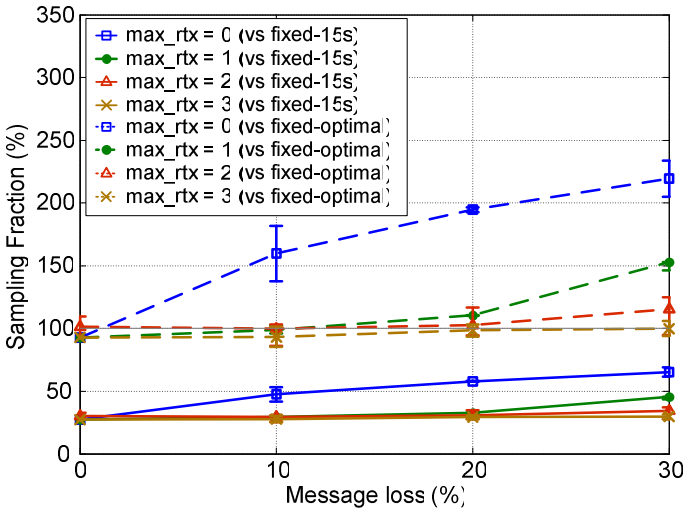

Figure 10. Sampling fraction for different max_rtx values.

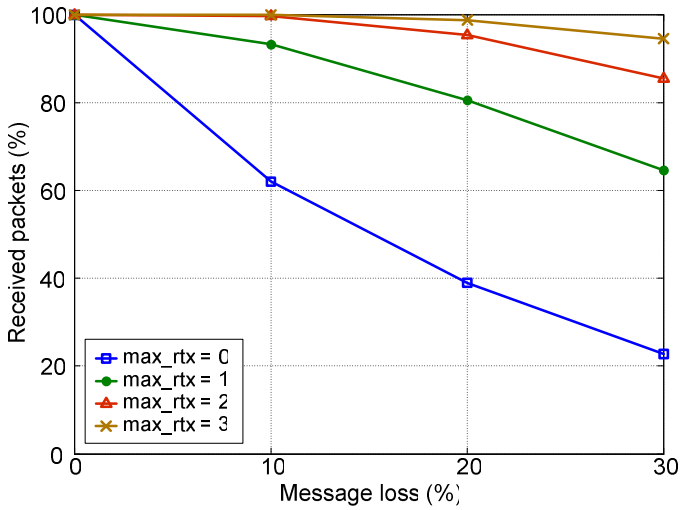

Figure 11. Delivery ratio for different *max_rtx* values.

In terms of energy consumption, a reduced Sampling Fraction caused by an increased delivery ratio immediately turns out into a lower sensor energy consumption, as shown in Figure 12. Things are not so straightforward for the energy consumption of the radio, which is given by the sum of two different components with contrasting behavior. On one hand, a large number of retransmissions leads to a high energy consumption. On the other, a low sampling frequency implies a lower number of messages to be transmitted.

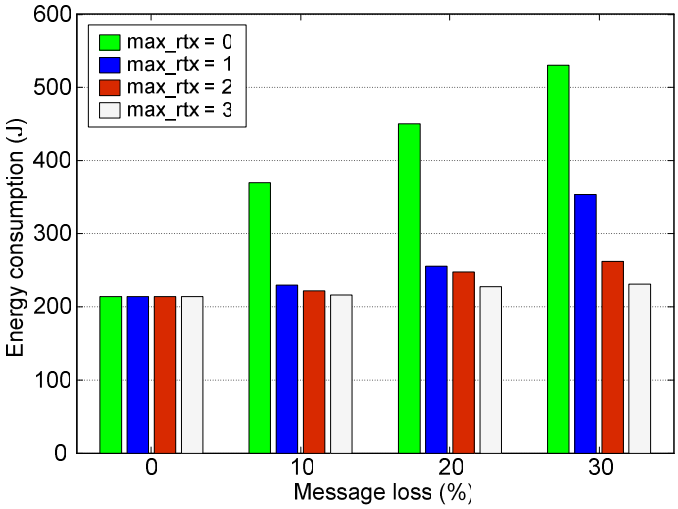

Figure 12. Total energy consumed by the sensor for different *max_rtx* values.

The total energy consumed by the radio equipment is shown in Figure 13. We can see that, in any case, it is much smaller than the energy consumed by the sensor. We appreciate the fact that ASA is really effective as the additional costs required for achieving a message delivery ratio

close to 100% are largely compensated by the reduction in the number of samples to be acquired and transmitted.

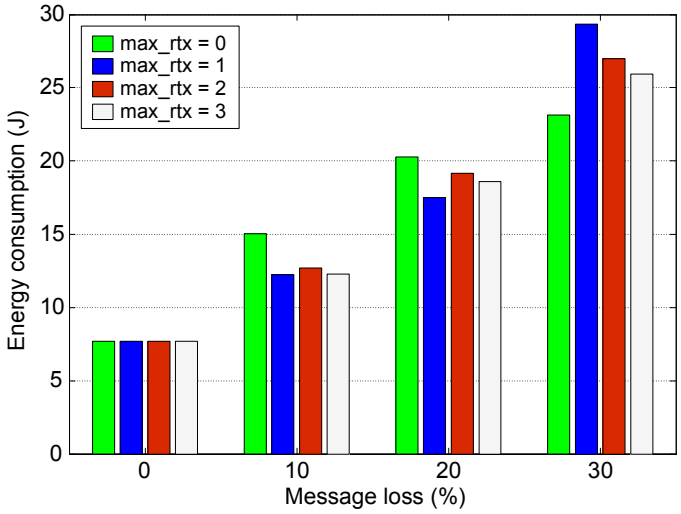

Figure 13. Total energy consumed by the radio for different *max_rtx* values.

Obviously, the above results strongly depend both on the specific sensor and the chosen sensor node platform. To show the effect of different sensor platforms, we also considered sensor nodes equipped with the CC2420 radio. The CC2420 radio is an evolution of the CC1000 one considered in previous experiments and is used, for example, in TmoteSky sensor nodes. It allows for a bit rate of 250 kbps (the bit rate provided by CC1000 is 19.2 kbps) and its power consumptions in transmit and receive modes are shown in Table 1 (here we assumed that the power consumption in idle mode is equal to that in receive mode, while the power consumption in the sleep mode is negligible). Figure 14 shows the total energy consumed by the sensor node for communication when using the two different radios, under the assumption that each message is re-transmitted up to two times. The energy consumption is significantly lower when using the CC2420 radio as the bit rate is more than one order of magnitude larger than the CC1000 one at the cost of a comparable power consumption. Since the energy cost for communication is lower, when using the CC2420, a larger number of retransmissions per message can be allowed. By increasing the delivery ratio, this results in a significant decrease in the sensor (i.e., overall) energy consumption.

The above results confirm the effectiveness of ASA in reducing the overall energy consumption in the presence of energy-hungry sensors. Moreover, the evolution from CC1000 to CC2420 is paradigmatic of a general trend observed in the last year in the field of wireless

technologies for sensor networks where there has been a significant increase in the bit rate provided by the sensor nodes with only a limited increase in their power consumption.

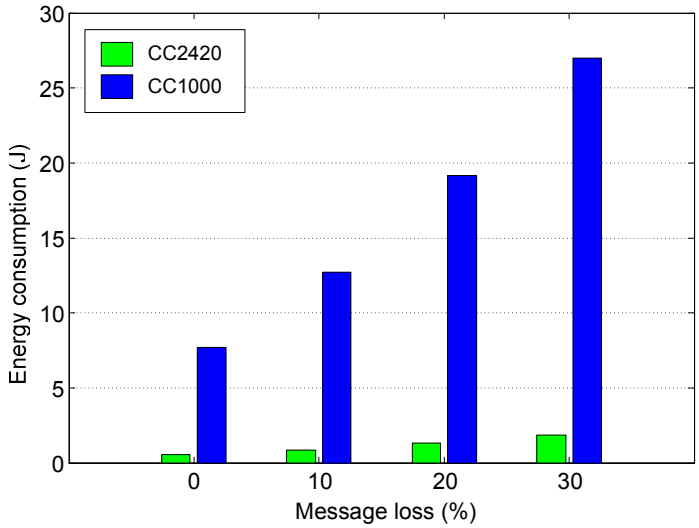

Figure 14. Comparison of total energy consumed for data communication when using the CC1000 and CC2420 radio equipments ($max\_rtx = 2$)

5. Conclusions

The paper proposes an Adaptive Sampling Algorithm (ASA) for wireless sensor networks which is capable of dynamically estimating the optimal sampling frequency of the acquired signals. The algorithm has been originally conceived to reduce the energy consumption of a prototype sensor for snow monitoring applications; however, the proposed approach is general and can be used in all cases where the process to be monitored exhibits a slow variation over time.

We performed an extended simulation analysis, based on traces derived from real measurements, by using the TOSSIM simulation tool. We found that ASA is able to reduce the number of acquired samples up to 79% w.r.t. the fixed sampling frequency, while generally preserving the accuracy of the data sequence collected at the base station. This results in a corresponding energy saving of both the sensor and the radio. We have also found that, thanks to its ability to adapt the sampling frequency to the real activity, our algorithm may perform similar to, or even better than, a fixed-rate scheme where the optimal sampling frequency is known in advance.

Since ASA exploits temporal correlation among successive data sample, it requires a message reliability close to 100% to work efficiently. We evaluated through simulation the cost, in term of additional energy consumption of the radio, for fulfilling this requirement. We found that benefits largely predominate over costs as the energy consumption of the overall system (i.e., both the sensor and the radio) is reduced. We are aware that this conclusion strongly depends on the specific sensor, whose power consumption is significantly larger than that of the radio but the analysis makes the point. In general, one should evaluate whether it is more convenient to acquire redundant data and tolerate some message loss, or minimize the number of acquired data and ask for a 100% reliability in message delivery. Obviously, the optimal strategy depends on the relative cost, in terms of energy consumption, for data acquisition the communication, i.e., on sensors and sensor nodes that are used.

Acknowledgments

This work is supported by the Italian Ministry for Education and Scientific Research (MIUR) under the FIRB ArtDeco. Authors wish to thank Dr. C. Galperti for having provided snow sensor data and Ing. F. Mancini for carrying out the simulation experiments.

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] A. Mainwaring, J. Polastre, R. Szewczyk, D. Culler and J. Anderson, "Wireless Sensor Networks for Habitat Monitoring", Proc. of the First ACM Workshop on Wireless Sensor Networks and Applications, Pages: 88-97 Atlanta, GA, USA, 28 September 2002.

[3] ON World Inc, "Wireless Sensor Networks – Growing Markets, Accelerating Demands", July 2005, available at http://www.onworld.com/html/wirelesssensorsrprt2.htm.

[4] Embedded WiSeNTs Consortium, "Embedded WiSeNts Research Roadmap (Deliverable 3.3)", available at http://www.embedded-wisents.org.

[5] C. Alippi and C. Galperti, "An adaptive system for optimal solar energy harvesting in wireless sensor network nodes," IEEE Transactions on Circuits and Systems I, Volume 55, Issue 6, July 2008, pp. 1742 - 1750.

[6] J. Polastre, "A Unifying Link Abstraction for Wireless Sensor Networks", Ph.D. dissertation, University of California at Berkeley, 2005. Available at http://www.polastre.com/papers/polastre-thesis-final.pdf.

[7] G. Anastasi, M. Conti, M. Di Francesco, A. Passarella, "Energy Conservation in Wireless Sensor Networks", Ad Hoc Networks, to appear (currently available at http://info.iet.unipi.it/~anastasi/papers/adhoc08.pdf).

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

[9] S.M. Mahmud, “High precision phase measurement using adaptive sampling”, IEEE Trans. Instrum. and Meas., vol. 38, issue 5, pp. 954-960, Oct. 1989.

[10] A. R. Varkonyi-Koczy , G. Simon , L. Sujbert and M. Fek “A fast filter-bank for adaptive Fourier analysis,” IEEE Trans. Instrum. Meas., vol. 47, pp. 1124, Oct. 1998.

[11] J.N. Lygouras, K.A. Lalakos and P.G. Tsalides “High-Performance Position Detection and Velocity Adaptive Measurement for Closed-Loop Position Control”, IEEE Trans. Instrum. and Meas., vol. 47, issue 4, pp. 978-985, Aug. 1998.

[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. L. Iman and W. J. Conover, A Modern Approach to Statistics. New York: Wiley, 1983.

[14] I. W. Burr, Statistical Quality Control Methods. New York: Dekker, 1976.

[15] R. Gnanadesikan,Methods for Statistical Data Analysis of Multivariate Observations. New York: Wiley, 1977.

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

[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] D.A, Rauth and V.T., Randal, "Analog-to-digital conversion. part 5," Instrumentation & Measurement Magazine, IEEE , vol.8, no.4, pp. 44-54, Oct. 2005

[19] Patent pending No. US 2006/0192568 A1.

[20] W. Heinzelman, A. Chandrakasan, H. Balakrishnan, “Energy-Efficient Communication Protocol for Wireless Microsensor Networks”, Proc. Hawaii International Conference on System Science (HICSS-33), January 2000.

[21] C. Alippi, G. Anastasi, M. Di Francesco, C. Galperti, F. Mancini, M. Roveri, “Effective Energy Management in Wireless Sensor Networks through Adaptive Sampling”, Technical report DII-TR-2008-08, Univ. of Pisa, Italy, 2008. Available at `http://info.iet.unipi.it/~anastasi/papers/DII-TR-2008-08.pdf`.

[22] P. Levis, N. Lee, M. Welsh, D. Culler, “TOSSIM Accurate and Scalable Simulation of Entire TinyOS Applications”, Proc. ACM SenSys 2003, Los Angeles, 2003.

[23] C. Alippi, G. Anastasi, C. Galperti, F. Mancini, M. Roveri, “Adaptive Sampling for Energy Conservation in Wireless Sensor Networks for Snow Monitoring Applications”, Proc. IEEE International Workshop on Mobile Ad hoc and Sensor Systems for Global and Homeland Security (MASS-GHS 2007), Pisa (Italy), October 8-12, 2007.

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

[25] S. Kim, R. Fonseca, D. Culler, “Reliable Transfer on Wireless Sensor Networks”, Proc. IEEE SECON 2004, pp. 449-459, 4-7 Oct., 2004.