An Adaptive Sampling Algorithm for Effective Energy Management in Wireless Sensor Networks With Energy-Hungry Sensors

Cesare Alippi, *Fellow, IEEE*, Giuseppe Anastasi, Mario Di Francesco, and Manuel Roveri

***Abstract*—Energy conservation techniques for wireless sensor networks generally assume that data acquisition and processing have energy consumption that is 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 that affect the network lifetime. In this paper, we propose an *adaptive sampling algorithm* that estimates online 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, 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.**

***Index Terms*—Adaptive systems, intelligent sensors, remote sensing, signal sampling, wireless local area network.**

I. Introduction

WIRELESS sensor networks (WSNs) are distributed measurement systems with a large number of measurement units that were 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, because 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 even more increase in the next years [3], [4]. However, energy consumption still remains a major obstacle for the full diffusion and exploitation of this technology, although batteries can be recharged, e.g., through solar energy harvesting mechanisms [5].

In recent 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 that is 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 aforementioned assumption does not necessarily hold. In fact, many real-life applications require specific sensors whose power consumption cannot be neglected [8]. Tables I and II provide the power consumptions of the most popular radio equipment 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 that aim at minimizing the radio activity need to be complemented with techniques that implement an efficient energy management of the sensors.

In this paper, we propose a general approach that leverages two complementary mechanisms at the sensor level: 1) *duty cycling* (i.e., the sensor board is switched off between two consecutive samples) and 2) *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 to measure the difference in phase between the fundamental components of two signals: 1) the sampling rate is increased until the phase is correctly measured or 2) the sampling rate reaches the maximum sampling rate of the system.

[10] describes a Fourier analyzer that autonomously adapts the parameters of the filters to match the signal components

Manuscript received August 27, 2008; revised March 13, 2009. First published August 18, 2009; current version published January 7, 2010. This work was supported in part by the Italian Ministry for Education and Scientific Research (MIUR) under the Fondo per gli Investimenti della Ricerca di Base Adaptive Infrastructures for Decentralized Organizations (FIRB ArtDeco). The Associate Editor coordinating the review process for this paper was Dr. Jesús Ureña.

C. Alippi and M. Roveri are with the Dipartimento di Elettronica e Informazione, Politecnico di Milano, 20133 Milano, Italy (e-mail: alippi@elet.polimi.it; roveri@elet.polimi.it).

G. Anastasi and M. Di Francesco are with the Department of Information Engineering, University of Pisa, 56122 Pisa, Italy (e-mail: giuseppe.anastasi@iet.unipi.it; mario.difrancesco@iet.unipi.it).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIM.2009.2023818

TABLE I
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 II
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

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 that 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 correspondingly adjust the sampling frequency.

ASA is more general than the aforementioned solutions, because it does not assume any hypothesis with regard to the nature of the signal (e.g., stationarity); moreover, its computational load is acceptable for midcomplexity WSN units.

ASA identifies the minimal sampling frequency online, which guarantees the reconstruction of the sampled signal; thus, 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 that will be transmitted and, as a consequence, the energy that the radio consumed. In addition, the proposed approach can be integrated with other techniques for energy conservation by acting at different abstraction levels (e.g., data aggregation and/or compression).

This paper is organized as follows. Section II introduces ASA. Section III presents the snow sensor that is used as a case study to assess the performance of the algorithm and the simulation environment that was used for performance analysis. Simulation results are finally presented in Section IV.

II. ASA: THE PROPOSED METHODOLOGY

The proposed ASA differentiates from the aforementioned 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 that was generally addressed with a statistical approach (for example, see [13]–[15]). Here, we found a modification of the cumulative sum (CUSUM) change detection test [16] to be particularly appropriate, which is widely used in the system-control community. In particular, we configured the test to detect changes that are associated with the highest frequency $F_{\max}$ of the signal, with $F_{\max}$ being related to the minimum sampling frequency F_N as per Nyquist $F_N > 2F_{\max}$ [17].

Frequency $F_{\max}$ is not available a priori and changes over time in a nonstationary process. Consequently, it clearly emerges that F_N also changes over time and that, by adapting the sampling frequency, oversampling is avoided, signal reconstruction is guaranteed, and power consumption reduction is 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, i.e., $\overline{F}_{\max}$ 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 to detect the change in the maximum frequency, we designed the following 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 a user-defined confidence

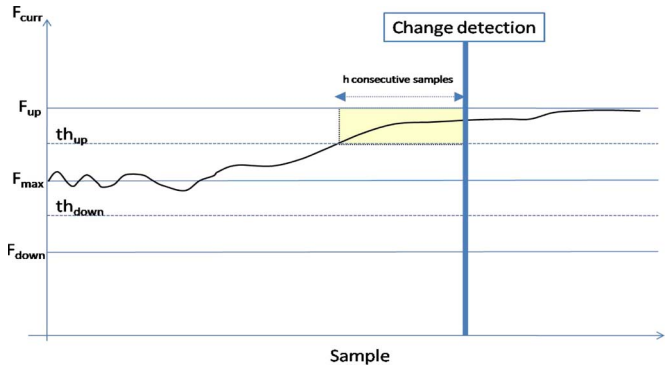

Fig. 1. Detecting a change in the maximum frequency.

parameter that 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 value is an independent parameter that the user/designer must set; for example, $\delta = 3\%$ implies that changes that affect $\overline{F}_{\max}$ for more than 3% $\overline{F}_{\max}$ must be detected (to be intended in statistical terms). F_{up} might significantly be 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. On the other hand, for higher δ, $F_{up} \cong (F_c/2)$, and the algorithm might suffer from the presence of false negatives. Obviously, F_{up} cannot be larger than $F_c/2$ due to the Nyquist theorem. F_{down} is less influenced by δ, because a decrement in the maximum frequency above $\delta\overline{F}_{\max}$ is detected (aliasing effects are not introduced here).

During the 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 following 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. One example of a frequency change is illustrated in Fig. 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 following detection rule: *if* $(|F_{curr} - F_{up}| < |F_{curr} - \overline{F}_{\max}|)$ *for* h *consecutive samples or if* $(|F_{curr} - F_{down}| < |F_{curr} - \overline{F}_{\max}|)$ *for* h *consecutive samples, then the new sampling frequency is* $F_c = cF_{curr}$.

The proposed algorithm is given in Algorithm 1.

Algorithm 1: ASA (c, δ, h)
1. Store the W initial samples that come from the process in the *Dataset;*
2. Estimate $\overline{F}_{\max}$ on the *Dataset* and set $F_c = c\overline{F}_{\max}$;
3. Define $F_{up} = \min\{(1+\delta)\cdot\overline{F}_{\max}, F_c/2\}$; $F_{down} = (1-\delta)\cdot\overline{F}_{\max}$;
4. $h_1 = 0$, and $h_2 = 0$; $i = W + 1$;
5. **while** (1){
6. Acquire the ith sample and add it to the *Dataset*;
7. Estimate the current maximum frequency F_{curr} on the sequence *Dataset* $(\mathrm{i} - W + 1, \mathrm{i})$;
8. **if** $(|F_{curr} - F_{up}| < |F_{curr} - \overline{F}_{\max}|)$
9. $h_1 = h_1 + 1$; $h_2 = 0$;
10. **else if** $(|F_{curr} - F_{down}| < |F_{curr} - \overline{F}_{\max}|)$ $h_2 = h_2 + 1$; $h_1 = 0$;
11. **else** $h_1 = 0$; $h_2 = 0$;
12. **if** $(h_1 > h) \| (h_2 > h)$ {
13. $F_c = cF_{curr}$;
14. $F_{up} = \min\{(1+\delta)\cdot\overline{F}_{\max}, F_c/2\}$;
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 to quickly detect a variation in the maximum frequency of the signal (but we could experience false positives that induce a continuous change in 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 for configuring h (T must be sufficiently large to grant that the estimate of h converges toward 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 $F_{up} = \min\{(1+\delta)\cdot\overline{F}_{\max}, F_c/2\}$; $F_{down} = (1-\delta)\cdot\overline{F}_{\max}$;
3. $h_1 = 0$, and $h_2 = 0$;
4. $\widetilde{h}_1 = 0$, and $\widetilde{h}_2 = 0$;
5. **for** $(\mathrm{i} = W + 1; \mathrm{i} < T; \mathrm{i}++)$ {
6. Estimate the current maximum frequency F_{curr} on sequence $(\mathrm{i} - W + 1, \mathrm{i})$
7. **if** $(|F_{curr} - F_{up}| < |F_{curr} - \overline{F}_{\max}|)$
8. $h_1 = h_1 + 1$; $h_2 = 0$;
9. **else if** $(|F_{curr} - F_{down}| < |F_{curr} - \overline{F}_{\max}|)$ $h_2 = h_2 + 1$; $h_1 = 0$;
10. **else** $h_1 = 0$; $h_2 = 0$;
11. **if** $(h_1 > \widetilde{h}_1)$ {
12. $\widetilde{h}_1 = h_1$;
13. }
14. **if** $(h_2 > \widetilde{h}_2)$ {
15. $\widetilde{h}_2 = h_2$;
16. }
17. }
18. **return** $(\min(\widetilde{h}_1, \widetilde{h}_2))$;

In more 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 (lines 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 $\widetilde{h}_2$). To be conservative, h is the minimum between $\widetilde{h}_1$ and $\widetilde{h}_2$ (line 18).

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

III. Experimental Setup

A. Description of the Snow Sensor

To evaluate the performance of ASA in a real application, we considered a sensor that was developed to monitor the snow composition (e.g., slope stability assessment and avalanches forecast). Such a sensor provides the dataset used in subsequent simulation experiments.

The snow sensor in this paper is a multifrequency capacitive measuring unit that was engineered to be embedded in a remote wireless measuring system. It is composed of a probe, a main multifrequency injection board that can measure capacity at different frequencies [19], and a wireless unit that will 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 proven 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 frequencies and the temperature information are 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 (e.g., calibration, electrode precharging, and charge sharing) in a cyclic way to obtain a reasonably stable reliable measure. This activity makes the sensor very energy consuming; for example, by sampling data every 15 s, the average energy that was consumed is 880 mJ/sample. Such a high value can be explained as follows: 1) the sensor is an ad hoc sensor that was not optimized for energy consumption and 2) 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 2 s. Such a choice leads to an energy consumption of approximately 150-mJ/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%).

B. Sensor Network Configuration

ASA can be implemented in any sensor network architecture. However, for simulation, we considered a cluster-based architecture (see Fig. 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 low-energy adaptive clustering hierarchy (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], which is a widely used simulator for WSNs. In the considered communication protocol, both nodes within a cluster and cluster heads use a time-division multiple-access scheme for exchanging data with the corresponding cluster head or base station, respectively. Each node (cluster head) remains active only during the time slots that were assigned for communication to minimize the radio energy consumption. Intracluster interferences (i.e., collisions due to simultaneous transmissions of nodes that belong to the same cluster) are thus avoided by the communication protocol, whereas intercluster interferences (i.e., collisions due to simultaneous transmissions by nodes that belong to different clusters) can still occur. We modeled the effects of possible intercluster interferences as message losses. Therefore, in our simulations, messages may be missed due to either transmission errors or intercluster interferences.

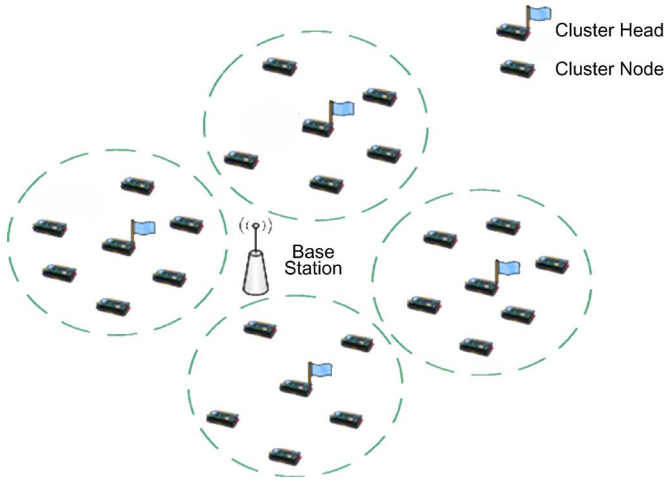

Fig. 2. Cluster-based sensor network architecture.

The communication protocol uses an automatic-repeat-request scheme based on acknowledgments, timeouts, and retransmissions to recover missed messages. Messages that were 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 that 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 step is a very simple approach that, nevertheless, proves to be effective in increasing the accuracy of the data sequence collected at the base station, although 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.

C. Figures of Merit

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

- *Sampling fraction*. This figure of merit is defined as the number of samples that were acquired by the sensor according to ASA with respect to the number of samples that were acquired with a fixed sampling frequency. The sampling fraction aims at evaluating the efficiency of the algorithm.
- *Sensor/Radio energy consumption*. This figure of merit summarizes the total energy that the sensor/radio subsystem consumed. The total energy consumption of the sensor is the product between the energy that each sampling cycle drained and the number of samples that were generated during the simulation experiment. The total energy that the radio consumed 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, and T_s represent the total time that the radio is in the transmitting, receiving, idle, or sleeping operational modes, whereas P_t, P_r, P_i, and 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 that was consumed in the sleeping mode is negligible with respect to the power that was consumed in the others modalities. Therefore, in our simulator, we implemented the approximated model as follows:

$$E_R = T_a \cdot P_a$$

  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)*. This figure of merit is defined as

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

  where x_i denotes the ith sample in the original data sequence, $\overline{x_i}$ denotes the ith sample in the data sequence reconstructed at the base station, and N denotes the total number of samples in the original data sequence, respectively. MRE gives a measure of the relative error that was 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 (MAE), which provides better indications than MRE in this specific case. It is defined as

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^{N} |\overline{x_i} - x_i|.$$

D. Parameter Settings and Methodology

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

TABLE III
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 that were derived from real measurements with the snow sensor described in Section III-A in different days and conditions. Each dataset consists of approximately 6.000 samples that were acquired with a fixed period of 15 s. This sampling frequency was chosen based on a priori knowledge of the signals that will be measured (e.g., snow capacitance and ambient temperature). It is large enough to capture quick variations (note that it is larger than necessary, because 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 a 90% confidence level [24].

In the following discussion, unless specified otherwise, figures refer to Experiment 2 being the most critical one for ASA. However, results are similar for other datasets.

IV. SIMULATION RESULTS

We divided our analysis into two parts. First, we investigated the advantages and disadvantages 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.

A. Adaptive Versus 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 (note that $F_c = c\overline{F}_{\max}$).

As presented in Fig. 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 that $F_{\max}$ and $F_c/2$ are the maximum frequency currently available in the signal and the maximum frequency detectable by ASA according to the Nyquist theorem, respectively. Obviously, when $F_c/2 < F_{\max}$, aliasing effects may occur, but ASA reacts by increasing the maximum sampling frequency.

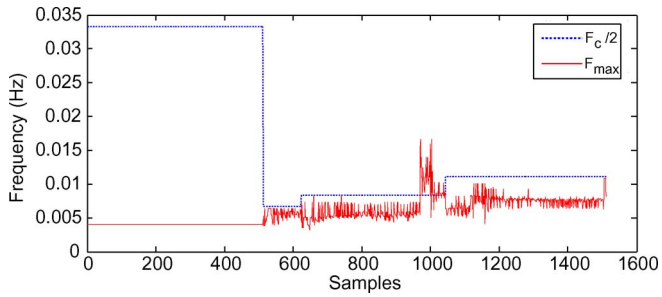

Fig. 3. Sampling rate as a function of the samples.

Initially, the sampling frequency was set to 1/15 Hz up to Sample 512 (we did not know the initial optimal sampling rate, and we opted to have it overdimensioned). Then, ASA reduces the sampling frequency to 1/75 Hz (obviously, any change in F_c is reflected on $F_c/2$); this approach allows us to reduce 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 1/60 Hz. Finally, the increase in F_{curr} around Sample 990 induces a next increment in the sampling frequency to 1/45 Hz. We note that the abrupt change in frequency that occurs at around Sample 990 introduces aliasing phenomena, because $F_{\max}$ is larger than $\overline{F}_{\max}$. Once that 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 min, because the sampling rate is about 1 sample/min, and $h = 40$).

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

We comment that, initially, the sensor node uses the maximum sampling rate 1/15 Hz. Afterward, 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 to adjust it to the current dynamics of the signals.

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

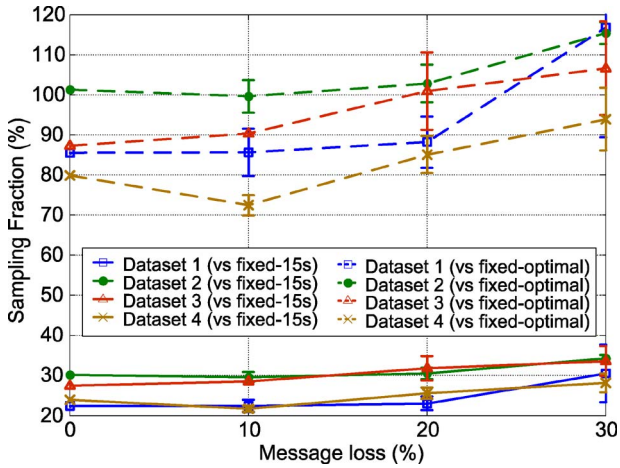

Fig. 4. Sampling fraction as a function of the message loss rate for different datasets.

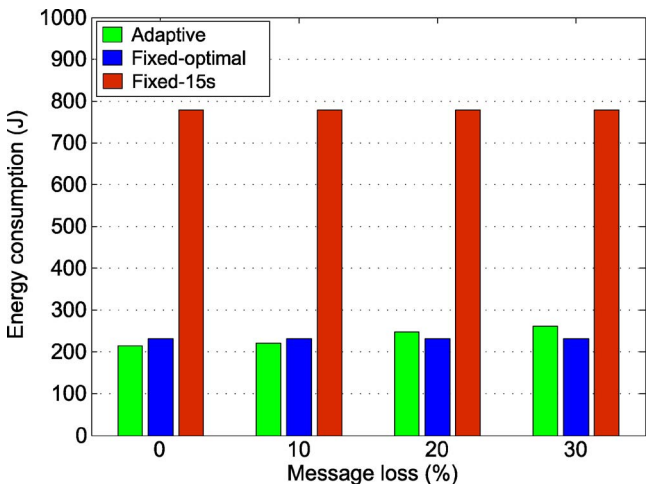

Fig. 5. Total energy that the snow sensor consumed.

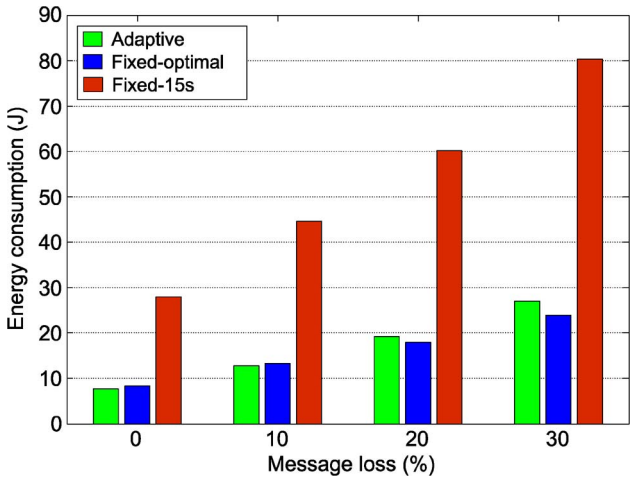

Fig. 6. Energy that the radio consumed.

The decrease in the sampling fraction that was provided by ASA results in a corresponding decrease in the energy consumption for both sensing and communication. Figs. 5 and 6 show the total energy that was consumed by the snow sensor and the radio, respectively, for Dataset 2. Note that the

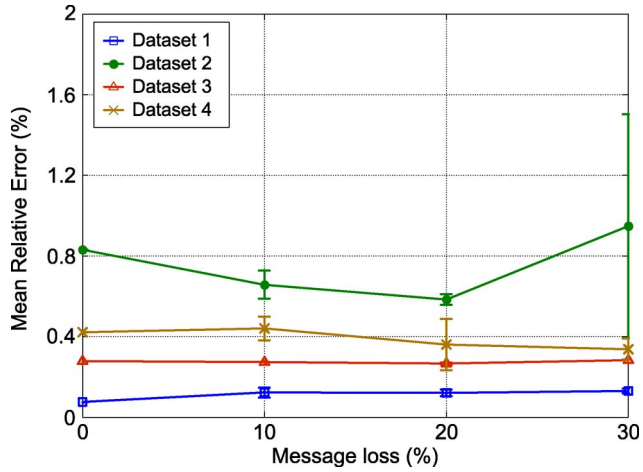

Fig. 7. MRE for low-frequency capacitance as a function of the message loss.

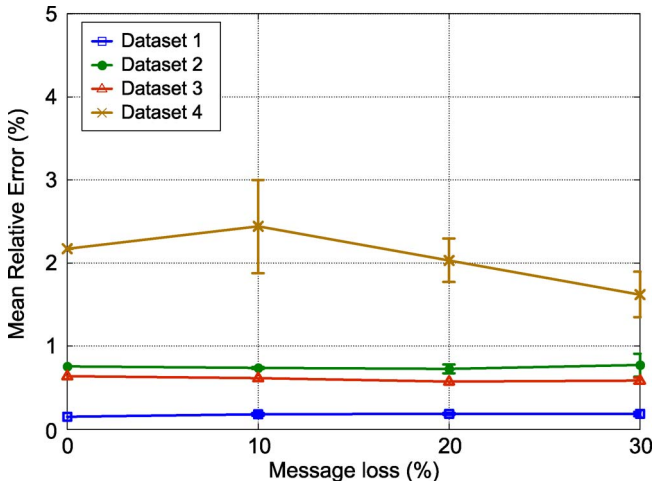

Fig. 8. MRE for high-frequency capacitance as a function of the message loss.

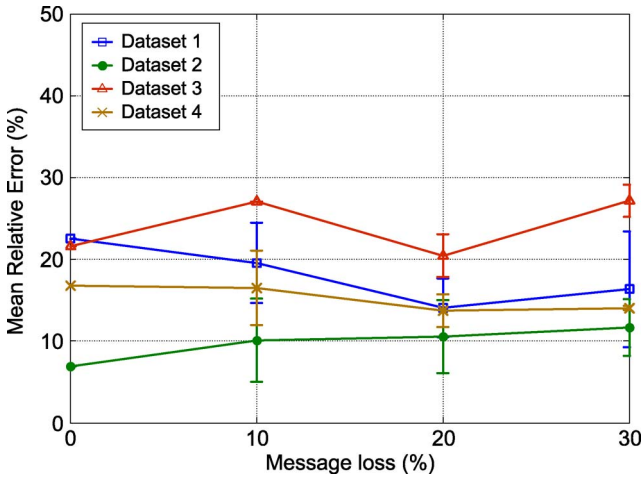

Fig. 9. MRE for temperature as a function of the message loss.

energy savings that were provided by ASA are not obtained at the expense of a decreased accuracy in the data sequence collected at the base station. Figs. 7 and 8 show the MRE for the snow capacitance at low and high frequencies, respectively, whereas Fig. 9 presents the same index for the temperature. The MRE, both at low and high frequencies, remains very low (i.e., 1%–2%), although 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 Fig. 9), 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 MAE, as summarized in Table IV, for the different datasets and message-loss probabilities. We can see that the average (absolute) deviation from the original temperature sequence is always negligible.

B. Impact of Communication Unreliability

According to the previous section, ASA can actually reduce the percentage of samples that will be acquired while assuring that the information will 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 Figs. 4–6), because ASA reduces the sampling frequency by exploiting the temporal correlation among consecutive samples. As such, to correctly work, ASA requires (almost) all data to be received by the base station. The algorithm can tolerate a certain fraction of missing samples due to the loss-compensation mechanism (the phenomenon under monitoring is assumed to slowly change 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 that were sent by the base station to sensor nodes to update 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 it is lower, aliasing effects may occur.

To make the adaptive sampling approach effective, we should thus guarantee a message delivery ratio (i.e., the percentage of messages correctly received by the final destination) in both directions close to 100%, e.g., through an acknowledgment-based retransmission protocol as considered here. In [25], it has been found that retransmission is the most efficient approach to data transfer reliability in WSNs. Obviously, message retransmission increases the delivery ratio but 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 range [0,3].

Fig. 10 shows the impact of the max_rtx value on the sampling fraction for increasing message-loss probabilities. As expected, the sampling fraction significantly decreases when the maximum number of retransmissions per message increases as the delivery ratio accordingly increases, as shown in Fig. 11. With $max\_rtx \geq 2$, more than the 85% of messages are delivered to the final destination (even with a link message-loss

TABLE IV
MAE FOR THE TEMPERATURE AS A FUNCTION OF THE MESSAGE LOSS (IN DEGREE CELSIUS)

	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

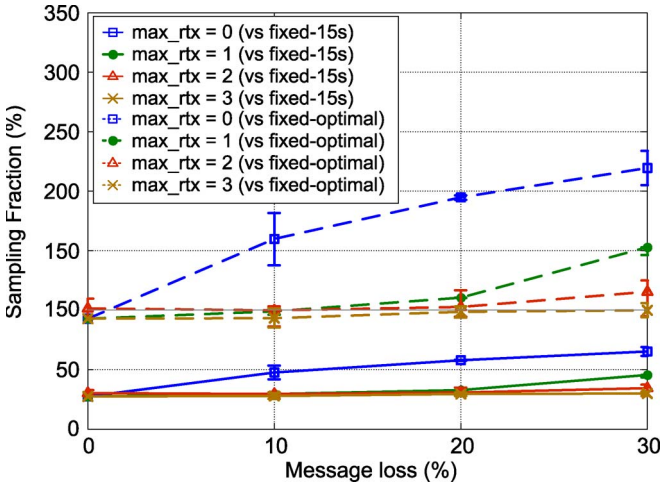

Fig. 10. Sampling fraction for different max_rtx values.

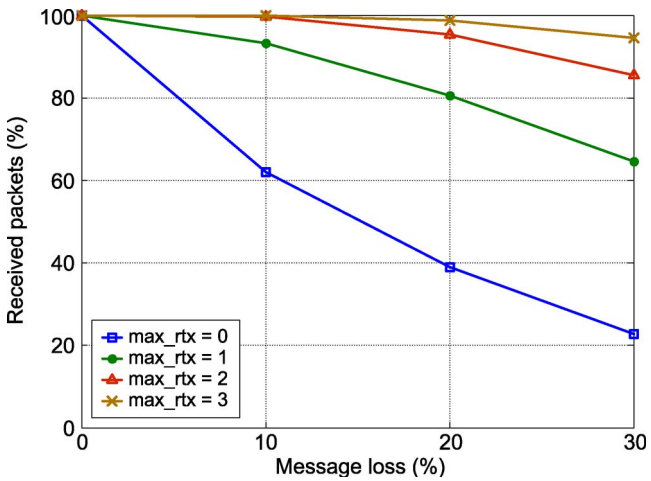

Fig. 11. Delivery ratio for different max_rtx values.

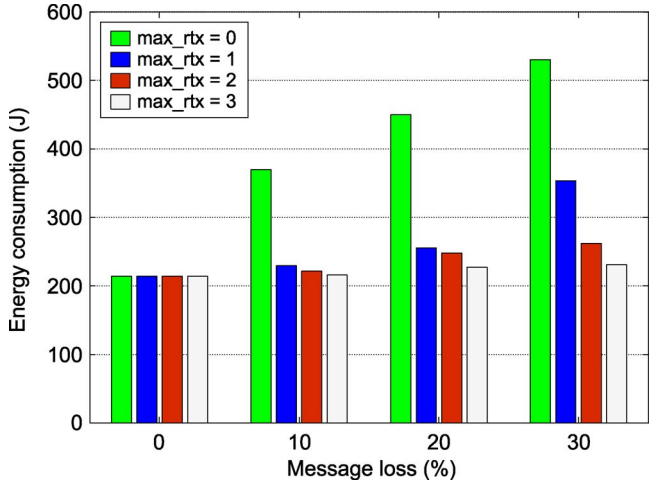

Fig. 12. Total energy that the sensor consumed for different max_rtx values.

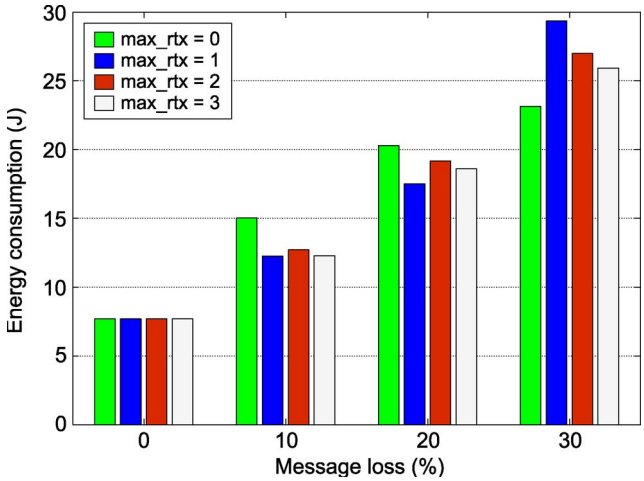

Fig. 13. Total energy that the radio consumed for different max_rtx values.

probability of 30%), and the performance of ASA is similar to or even better than that of the (unfeasible) fixed-rate approach.

In terms of energy consumption, a reduced sampling fraction that was caused by an increased delivery ratio immediately turns out into a lower sensor energy consumption, as shown in Fig. 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 lead to a high energy consumption. On the other hand, a low sampling frequency implies a lower number of messages that will be transmitted.

The total energy that the radio equipment consumed is shown in Fig. 13. We can see that, in any case, it is much smaller than the energy that the sensor consumed. We appreciate the fact that ASA is really effective as the additional costs that were required to achieve a message delivery ratio that is close to 100% are largely compensated by the reduction in the number of samples that will be acquired and transmitted.

Obviously, the aforementioned 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 that were equipped with the CC2420 radio. The CC2420 radio is an evolution of the CC1000 one, which was considered in previous experiments and is used, for example, in TmoteSky sensor nodes. It allows for a bit rate of 250 kb/s (the bit rate that was provided by CC1000 is 19.2 kb/s), and its power consumptions in the transmit and receive modes are shown in Table I (here, we assumed that the power consumption in the idle mode is equal to that in the receive mode, whereas the power consumption in the sleep mode is

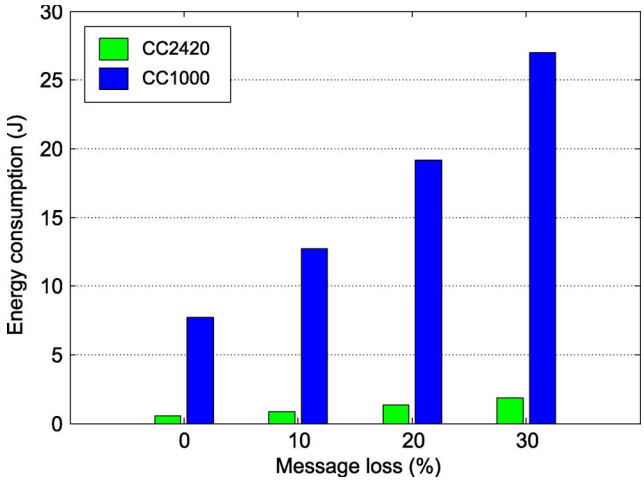

Fig. 14. Comparison of the total energy that was consumed for data communication when using the CC1000 and CC2420 radio equipment ($max\_rtx = 2$).

negligible). Fig. 14 shows the total energy that was consumed by the sensor node for communication when using the two different radios, under the assumption that each message is retransmitted up to two times. The energy consumption is significantly lower when using the CC2420 radio, because the bit rate is more than one order of magnitude larger than the CC1000 one at the cost of a comparable power consumption. The energy cost for communication is lower; thus, when using the CC2420, a larger number of retransmissions per message can be allowed. By increasing the delivery ratio, this approach results in a significant decrease in the sensor (i.e., overall) energy consumption.

The aforementioned 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 that has been observed in recent years in 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.

V. Conclusion

This paper has proposed an ASA for WSNs, which can dynamically estimate the optimal sampling frequency of the acquired signals. The algorithm has originally been 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 that will be monitored exhibits a slow variation over time.

We performed an extended simulation analysis based on traces from real measurements by using the TOSSIM simulation tool. We found that ASA can reduce the number of acquired samples up to 79% with respect to the fixed sampling frequency while generally preserving the accuracy of the data sequence collected at the base station. This approach results in a corresponding energy saving of both the sensor and the radio. We have also found that, due 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.

ASA exploits temporal correlation among successive data sample; thus, it requires a message reliability close to 100% to efficiently work. Through simulation, we evaluated the cost, in terms of additional energy consumption of the radio, to fulfill 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 of the communication, i.e., on sensors and sensor nodes that were used.

Acknowledgment

The authors would to thank Dr. C. Galperti for providing snow-sensor data and Ing. F. Mancini for carrying out the simulation experiments.

References

[1] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "Wireless sensor networks: A survey," *Comput. Netw.*, vol. 38, no. 4, pp. 393–422, Mar. 2002.

[2] A. Mainwaring, J. Polastre, R. Szewczyk, D. Culler, and J. Anderson, "Wireless sensor networks for habitat monitoring," in *Proc. 1st ACM Workshop Wireless Sensor Netw. Appl.*, Atlanta, GA, Sep. 28, 2002, pp. 88–97.

[3] ON World Inc, *Wireless Sensor Networks: Growing Markets, Accelerating Demands*, Jul. 2005. [Online]. Available: http://www.onworld.com/html/wirelesssensorsrprt2.htm

[4] Embedded WiSeNTs Consortium, *Embedded WiSeNts Research Roadmap (Deliverable 3.3)*. [Online]. Available: 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 Trans. Circuits Syst. I: Reg. Papers*, vol. 55, no. 6, pp. 1742–1750, Jul. 2008.

[6] J. Polastre, "A unifying link abstraction for wireless sensor networks," Ph.D. dissertation, Univ. of California Press, Berkeley, CA, 2005.

[7] G. Anastasi, M. Conti, M. Di Francesco, and A. Passarella, "Energy conservation in wireless sensor networks," *Ad Hoc Networks*, vol. 7, no. 3, pp. 537–568, May 2009.

[8] V. Raghunathan, S. Ganeriwal, and M. Srivastava, "Emerging techniques for long-lived wireless sensor networks," *IEEE Commun. Mag.*, vol. 44, no. 4, pp. 108–114, Apr. 2006.

[9] S. M. Mahmud, "High-precision phase measurement using adaptive sampling," *IEEE Trans. Instrum. Meas.*, vol. 38, no. 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, no. 5, pp. 1124–1128, 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. Meas.*, vol. 47, no. 4, pp. 978–985, Aug. 1998.

[12] A. Jain and E. Y. Chang, "Adaptive sampling for sensor networks," in *Proc. Workshop DMSN*, Toronto, ON, Canada, 2004, pp. 10–16.

[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: Marcel 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*. Englewood Cliffs, NJ: Prentice-Hall, 1993.
[17] A. J. Jerri, "The Shannon sampling theorem—Its various extensions and applications: A tutorial review," *Proc. IEEE*, vol. 65, no. 11, pp. 1565–1595, Nov. 1977.
[18] D. A. Rauth and V. T. Randal, "Analog-to-digital conversion part 5," *IEEE Instrum. Meas. Mag.*, vol. 8, no. 4, pp. 44–54, Oct. 2005.
[19] E. Pasero, M. Riccardi, and T. B. Meindl, "Multi-frequency capacitive measurement device and a method of operating the same," U.S. Patent App. 11/228 065, Sep. 16, 2005.
[20] W. Heinzelman, A. Chandrakasan, and H. Balakrishnan, "Energy-efficient communication protocol for wireless microsensor networks," in *Proc. 33rd HICSS*, Jan. 2000, pp. 1–10.
[21] C. Alippi, G. Anastasi, M. Di Francesco, C. Galperti, F. Mancini, and M. Roveri, "Effective energy management in wireless sensor networks through adaptive sampling," Univ. Pisa, Pisa, Italy, Tech. Rep. DII-TR-2008-08, 2008.
[22] P. Levis, N. Lee, M. Welsh, and D. Culler, "TOSSIM: Accurate and scalable simulation of entire TinyOS applications," in *Proc. ACM SenSys*, Los Angeles, CA, 2003.
[23] C. Alippi, G. Anastasi, C. Galperti, F. Mancini, and M. Roveri, "Adaptive sampling for energy conservation in wireless sensor networks for snow monitoring applications," in *Proc. IEEE Int. Workshop MASS-GHS*, Pisa, Italy, Oct. 8–12, 2007, pp. 1–6.
[24] J. Banks, *Handbook of Simulation*. New York: Wiley, 1998.
[25] S. Kim, R. Fonseca, and D. Culler, "Reliable transfer on wireless sensor networks," in *Proc. IEEE SECON*, Oct. 4–7, 2004, pp. 449–459.

Cesare Alippi (SM'94–F'06) received the Dr.Ing. degree *(summa cum laude)* in electronic engineering and the Ph.D. degree in computer engineering from the Politecnico di Milano, Milan, Italy, in 1990 and 1995, respectively. He has completed his research work in computer sciences at the University College, London, U.K., and the Massachusetts Institute of Technology, Cambridge.

He is currently a Full Professor of information processing systems with the Dipartimento di Elettronica e Informazione, Politecnico di Milano. His research interests include application-level analysis and synthesis methodologies for embedded systems, computational intelligence, and wireless sensor networks. His research results have been published in more than 120 technical papers in international journals and conference proceedings.

Dr. Alippi is an Associate Editor for the IEEE TRANSACTIONS ON NEURAL NETWORKS and the IEEE TRANSACTIONS ON INSTRUMENTATION AND MEASUREMENT.

Giuseppe Anastasi received the M.S. degree in electrical engineering and the Ph.D. degree in computer engineering from the University of Pisa, Pisa, Italy, in 1990 and 1995, respectively.

He is currently an Associate Professor of computer engineering with the Department of Information Engineering, University of Pisa. He is an Area Editor of the Elsevier journals *Pervasive and Mobile Computing* and *Computer Communications*. He is a Coeditor of the book *Advanced Lectures in Networking* (LNCS 2497, Springer, 2002) and has published more than 70 papers in computer networking and pervasive computing. His research interests include pervasive computing, sensor networks, and power management.

Dr. Anastasi is a member of the IEEE Computer Society. He was a Program Cochair of the Ninth IEEE International Symposium on a World of Wireless, Mobile, and Multimedia Networks (WoWMoM 2008), the Vice Program Chair of the Fourth IEEE International Conference on Mobile Ad Hoc and Sensor Systems (MASS 2007), a General Cochair of the IEEE WoWMoM 2005, the Workshops Chair of the Fourth IEEE Conference on Pervasive Computing and Communications Workshops (PerCom 2006), the IEEE WoWMoM 2006, and the 16th IEEE International Conference on Computer Communications and Networks (ICCCN 2007).

Mario Di Francesco received the B.S. and M.S. degrees in computer engineering from the University of Pisa, Pisa, Italy, in 2002 and 2005, respectively. He is currently working toward the Ph.D. degree in the Department of Information Engineering, University of Pisa.

His research interests include pervasive computing and wireless sensor networks.

Mr. Di Francesco was a Publications Cochair of the Seventh IEEE International Symposium on a World of Wireless, Mobile, and Multimedia Networks (WoWMoM 2006) and currently serves on the Technical Program Committee of the Fifth IEEE International Workshop on Sensor Networks and Systems for Pervasive Computing (PerSeNS 2009)

Manuel Roveri received the Dr.Eng. degree in computer science engineering from the Politecnico di Milano, Milan, Italy, in June 2003, the M.S. degree in computer science from the University of Illinois at Chicago, Chicago, in December 2003, and the Ph.D. degree in computer engineering from the Politecnico di Milano in May 2007.

He is currently an Assistant Professor with the Dipartimento di Elettronica e Informazione, Politecnico di Milano. His research interests include computational intelligence, adaptive algorithms, and wireless sensor networks.