Predictive models for energy saving in wireless sensor networks

Article

Accepted version

A. De Paola, G. Lo Re, F. Milazzo, M. Ortolani

In Proceedings of the IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2011, pp. 1-6

Publisher: IEEE

http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5986204

Predictive Models for Energy Saving in Wireless Sensor Networks

Alessandra De Paola, Giuseppe Lo Re, Fabrizio Milazzo, and Marco Ortolani
Department of Computer Engineering – University of Palermo
Viale delle Scienze, ed 6. – 90128 Palermo, Italy
Email: {depaola, lore, fmilazzo, ortolani}@unipa.it

***Abstract*—ICT devices nowadays cannot disregard optimizations toward energy sustainability. Wireless Sensor Networks, in particular, are a representative class of a technology where special care must be given to energy saving, due to the typical scarcity and non-renewability of their energy sources, in order to enhance network lifetime. In our work we propose a novel approach that aims to adaptively control the sampling rate of wireless sensor nodes using prediction models, so that environmental phenomena can be consistently modeled while reducing the required amount of transmissions; the approach is tested on data available from a public dataset.**

***Keywords*-Wireless Sensor Networks; Predictive Models; Energy Saving.**

I. Introduction and Motivations

In the past few years, the increasing social awareness about the negative impact on the environment deriving from the mere use of ICT equipment has boosted the research on methodologies and algorithms for optimizing the energy consumption of the involved devices. For instance, as mentioned in [1], the Internet alone is responsible for a consumption amounting to $\sim$74 TeraWatts hours (TWh) per year, only in the USA. This figure might be significantly reduced by careful application of power management methods, but the issue is seldom properly addressed during the design phase of commonly available systems. Both the potential repercussions to the surrounding environment, and the lack of proper use of the available resources are thus likely disregarded. The goal of any energy sustainable system is thus twofold, and ought to encompass both an *outer* optimization (considering the global expected outcome), and *inner* optimization (regarding the optimal utilization of the basic system components).

Both aspects are especially emphasized whenever resource-constrained devices are involved, as is the case, for instance, of Wireless Sensor Networks (WSNs), which we will specifically consider in this paper. In the past few years, WSN technology has grown into one of the most promising tools for collecting data and extracting information from remote or hostile sites [2]. Wireless sensors nodes are comparable to fully functional computers, in that they are not just able to collect measurements of physical quantities, but also to perform limited computations; the most remarkable difference between them and traditional computing devices is represented by their limited power supply (typically, batteries).

An interesting usage scenario for WSNs regards the design of energy efficient buildings [3], [4], where they represent the perceptive component of a complex system aimed at reducing the overall energy impact, and represents a typical use of such tool for outer optimization. At the same time, an effective usage of such technology cannot disregard the inner optimization of the energy sources available to each of the sensor nodes; for instance, even when renewable energy sources are available, they may not be constantly accessible (e.g. solar cells at nighttime), so effective resource allocation, such as adaptive load balancing, must be enforced [5].

Our work specifically aims at extending the network lifetime by enforcing careful usage of the energy resources available to nodes. In [6] the main sources of energy consumption in wireless sensor nodes are identified in the sensing, processing, and communication components; moreover, it has been shown that the transceiver represents the major drain, especially as compared to the CPU [7]. Besides implementing efficient MAC protocols including duty cycle techniques for turning on the radio only when strictly necessary [8], energy efficiency may thus be achieved by exploiting the inherent spatial and temporal redundancy in data. Such physical quantities as temperature, humidity, and pressure typically exhibit smooth variations, and change slowly over time [9], so that the node computational capabilities may be profitably used to extract predictive models for avoiding unnecessary sensing.

The main idea of our work consists in exploiting the temporal redundancy of measured physical quantities in order to compute predictive models, allowing to adaptively set the sensing rate of the nodes, and consequently reduce the overall amount of required transmissions. Our method has been tested on real measurements obtained from a publicly available repository containing data collected in an indoor environment, for which our underlying hypotheses hold. We assume that the wireless sensor network is arranged according to a hierarchical, cluster-based structure, where nodes forward sensed data towards their representative cluster head, which models the trend of the physical quantities by performing an on-line fitting through a mixture of Gaussians; models are kept up-to-date by ensuring that the error falls below a pre-defined threshold. We prove the energy efficiency of the proposed method by showing that the overall number of required sensory readings may be reduced with respect to a non adaptive sampling, without negatively affecting reliability.

II. Related Works

Optimizing energy consumption in wireless sensor networks is a widely studied issue, and the authors of [10] propose a taxonomy of various approaches presented in recent literature, dividing them into three main categories: *duty cycling*, *mobility-based*, and *data-driven*.

Duty cycling approaches specifically target the optimization of the networking subsystem, mainly focusing on the implementation of efficient algorithms for controlling the sleep/wakeup schedule of the radio transceiver; such approaches are typically not specifically concerned with the data sensed by nodes.

The second category is represented by mobility-based approaches, which allow for the implementation of higher-level techniques, such as load balancing, data muling, or energy harvesting; however, they impose strict requirements on the needed hardware.

Finally, data-driven approaches are more tightly bound to the intrinsic nature of sensed data; they often rely on the predictability of the monitored physical quantities, which are thus reliably representable through mathematical models. Such approaches are particularly relevant to the topic discussed in this paper.

The authors of [11] propose a method exploiting the correlation of data sensed by close nodes in order to build predictive models. In their system, nearby nodes are grouped into clusters, and it is assumed that the cluster head acts as a representative for all nodes within its group. The intrinsic spatial correlation of data allows some of the nodes to go into a *sleep* state, while the cluster head keeps on sensing; the role of representative is routinely taken up by all nodes within the group. During its turn, each cluster head also computes a predictive model for data, which will be shared within its cluster, and used instead of actual sensing as long as it is deemed reliable. Such technique results in an overall reduction of the required transmissions.

In [12], an autoregressive model is built using sensor readings. Initially, readings are collected until a buffer is filled; successively, each node computes a model for the sensed data and transmits only the model parameters back to the base station, thus effectively implementing a compression of data. The model is constantly checked for reliability against new readings; if a sufficiently large number of readings is recognized to fall behind a tolerance threshold within a given time window, the model is invalidated and recomputed by filling the buffer with fresh readings. However, if the readings are recognized as outliers, they are simply discarded, and the model will still be valid.

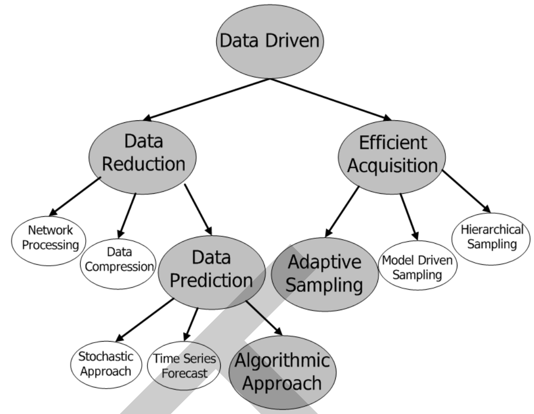

Figure 1. Categorization of data-driven approaches, according to [10]; shadowed ellipses highlight the methods exploited in our work.

Finally, the authors of [13] suggest to avoid the exchange of models through the network, by computing two separate predictors on the source and the sink node. Synchronization of such models is obtained by a minimal information exchange, consisting in the set of readings not satisfying a user-defined threshold, that signals the unreliability of the model computed at the source.

Unlike the previous methods, our proposal completely avoids the need of building a prediction model on sensor nodes, and delegates the whole computational burden to the cluster head. Sensor nodes are just required to sense the environmental phenomena with an adaptive sampling rate, thus reducing both the the number of necessary computations and of transmissions toward the cluster head. Reliability of the computed model is checked whenever new data are received by the cluster head, which will set the sampling rate of the source nodes accordingly (by decreasing it, if the model is reliable, and increasing it, otherwise).

III. Prediction Models for Energy Saving

This section describes the method devised for reducing the energy consumption of a wireless sensor network for environmental monitoring by exploiting the temporal redundancy inherent to the measured physical quantities. According to the taxonomy proposed in [10], the main contributions of our work may be classified as data reduction through prediction with algorithmic approach, and efficient data acquisition through adaptive sampling as shown in Figure 1.

Typical sensor nodes measure physical quantities which depend on several and possibly complex environmental interactions, and may be represented through analytical or empirical models. The former capture the result of the interactions by considering a large number of environmental variables, but their main drawbacks are the difficulty of generalization for a different physical phenomena, and their mathematical complexity which negatively impacts their computational viability. Alternatively, an empirical formulation may be given in order to model the considered physical

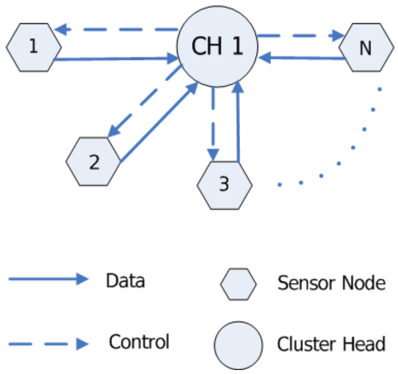

Figure 2. Example of a star topology cluster, highlighting the nature of the communications between the cluster head and each of the nodes.

phenomenon. Unlike the previous case, physical quantities are modeled by analyzing mathematical properties of their measured trends (e.g. periodicity, correlation, and so on) over a given amount of time, rather than studying their physical nature in order to predict future values.

In our work, we assume that the sensor network is organized according to a cluster-based topology, as depicted by Figure 2, where leaf nodes forward sensed data towards their representative cluster head, which in turn computes a model for the considered physical quantity, and sets the sampling rate for each of its children nodes in order to keep the model reliable over time. Moreover, we assume that the role of cluster head is assigned to more powerful wireless sensor nodes, the so-called *micro-servers* (e.g. Stargate nodes), which are typically equipped with larger amounts of memory, and offer more computational power. Given the characteristics of the considered quantities, we expect that some kind of periodicity is loosely present in data, so we only assume that the overall "shape" of the function representing a physical quantity is preserved over time.

Let $f(t)$ represent the trend of any observed physical quantity, defined over a time period Δt. Such function can be derived by data fitting, which in our case was approximated by a Gaussian mixture, as described by the following equation:

$$f(t) = \sum_{j=0}^{M-1} w_j \mathcal{N}(t|\mu_j, \sigma_j^2), \tag{1}$$

where mean μ_j, variance σ_j, and weight w_j are computed for each of the M Gaussian functions, through the Expectation Maximization algorithm [14].

According to the natural periodicity of the physical phenomena, we assume that the observed data are highly autocorrelated; the trend of $f(t)$ predicted for time interval Δt^{new} will thus be correlated to the one observed over a past time interval Δt^{old}. This can be expressed through a general shape-preserving transformation, given by the equation:

$$f^{new}(t) = \alpha \left[f^{old}(t) - \theta \right] + \theta, \tag{2}$$

where α and θ represent scaling and translation parameters of the transformation, respectively.

Parameters α and θ can be computed for a new time interval Δt_{new} by considering a new observation set S^{new}, and a model $f^{old}(t)$ for a past time interval Δt_{old}, through least squares optimization:

$$\underset{\alpha,\theta}{\mathbf{argmin}} \sum_{t \in \Delta t^{new}} [f^{new}(\alpha, \theta, t) - S^{new}(t)]^2. \tag{3}$$

The entire model is therefore computed and updated on the cluster head through an iterative algorithm that accepts as input a fitting model $f^{old}(t)$ for the observed set in a past time interval Δt^{old}, and computes the new model $f^{new}(t)$ in an on-line fashion. The raw samples gathered from the sensor nodes are iteratively processed at the cluster head to refine the estimated values of α and θ.

Moreover, in order to reduce the overall energy consumption, the algorithm adaptively controls the sampling rate of the sensor node by decreasing it when the absolute difference between prediction and actual measurement is lower than a given threshold, and increasing it otherwise. Such behavior makes the sampling process faster when the estimated model is not sufficiently reliable; on the other hand, sensor nodes are allowed to save energy (relative to sensing, and network transmissions) because the sampling period is reduced. Intuitively, the fact that the model becomes less reliable may be due to the presence of sudden bursts of high-frequency data (noise, or sudden variations) which should trigger a greater accuracy for the model itself.

The iterative algorithm starts by initializing all the necessary variables:

$$\theta \leftarrow 0, \quad \alpha \leftarrow 1, \quad T \leftarrow 1, \quad R \leftarrow \emptyset,$$

where T is the factor controlling the sampling rate, and R is the set of the current sensory readings.

The core of the algorithm consists of a read-prediction loop; at each step, the value of the prediction at time t is computed by the cluster head according to Equation 2 by using the current values of α and θ. Meanwhile, sensor nodes in the same cluster keep collecting samples of the environmental quantities at their current rate, and send them toward the cluster head, where they are appended to the set R of the current sensory readings, causing the deletion of the oldest ones. The accuracy of the prediction model is tested by comparing the latest measurement with the corresponding prediction. If the absolute difference is lower than a pre-defined threshold, the sampling rate is increased by doubling the controlling factor ($T \leftarrow 2 \cdot T$); otherwise, the model needs to be adapted to potentially fast variations in the monitored physical phenomenon, so the sampling rate is reset ($T \leftarrow 1$).

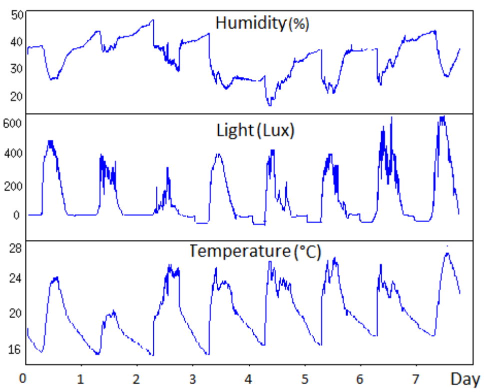

Figure 3. The plots show the average trends for temperature, humidity, and light for a representative set of nodes of the Intel Berkeley Research Lab; the rough periodicity of such patterns is clearly visible.

Regardless of the threshold, whenever a sensory reading is received by the cluster head, the transformation parameters are updated via a gradient descent algorithm starting from the latest estimates, according to Equation 3.

IV. Case Study and Experimental Results

In order to assess the approach described in the previous Section, we carried out a set of experiments using real data collected by Mica2-dot sensor nodes, and contained in the public database available from [15]. The reference environment is the Intel Berkeley Research Lab, containing 54 sensor nodes able to sense typical physical quantities: temperature, relative humidity, and light exposure. The measured values range from 17 to 25°C for temperature, from 20 to 90% for relative humidity, and from 0 to 1800 Lux for light exposure; the observations were carried out over a time period spanning several days, between February, 28 2004 and April, 5 2004; the sampling rate of each node was about 31s.

In this work we considered a time interval of nine days, from February, 29 2004 to March, 8 2004, and preprocessed the database in order to get rid of incomplete or clearly erroneous measurements (quasi-dead battery being a conspicuous source of error). We preliminary observed that all of these physical quantities show periodical patterns, as is evident by looking at Figure 3, which represents the underlying hypothesis for our work. The validity of such consideration has been experimentally confirmed by checking that the monitored quantities did present intrinsic redundancy, as shown by the autocorrelation plots reported in our previous work [16]; namely, both temperature and humidity present autocorrelation higher than 0.9, while for light is is higher than 0.75.

The size of the Intel Berkeley Research Lab is about $30 \times 40m^2$. In order to assess the precision of our prediction model, we need to consider clusters of nearby nodes, whose measurements are temporally correlated. An accurate study

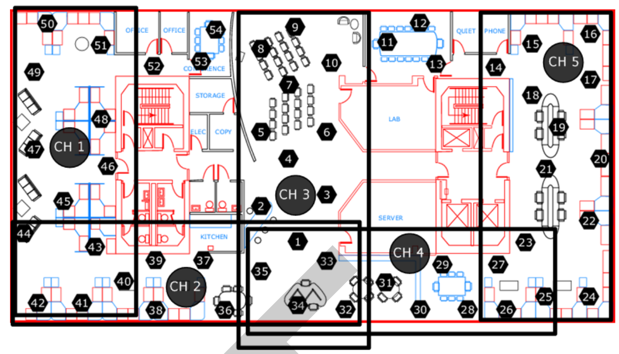

Figure 4. Location of the sensor nodes deployed at the Intel Berkeley Research Lab [15]; rectangles highlight the areas identified for our experiments.

reported in [17] identified five areas, where the measured quantities show high spatial correlation, so we chose to use the same layout for the experiments in this work. In the following we verify the validity of our assumptions, and show that the application of our algorithm has good potential for energy saving, in the hypothetical scenario, where each area is managed by a cluster head directly connected to each of the original nodes, as shown in Figure 4.

A. Predictive models validation

In order to obtain uniform comparison for our experimental results, we chose to normalize our metrics (mean, variance and maximum absolute prediction error), with respect to the previously mentioned ranges of each physical quantity. Each of the cluster heads depicted in Figure 4 is assumed to run the previously described algorithm for computing a prediction model for its nodes; the threshold used to invalidate the model depending on the tolerable prediction error was set to 5% for our experiments, and the periodicity of each physical quantity was assumed to be 24 hours.

The reliability of the produced predictive models has been assessed by computing the mean absolute error for each area and for each physical quantity. Figure 5 plots the the mean absolute prediction error for the test area 1 for temperature, humidity, and light, respectively; other areas show analogous trends. The values of the mean absolute error for all areas are reported in Table I, while Table II reports the variance and Table III the maximum. It is worth noting that the mean prediction error is very low in all cases (about 1%); in addition the largest variance is observed for light in area 4, but it is still not very high (2.25%); moreover, peaks of error only occur occasionally and remain localized to very few time instants, with the maximum amounting to ∼6% for light in areas 1 and 4.

B. Energy saving assessment

The approach described in Section III, and the cluster-based topology that we assume superimposed on the considered sensor field allow us to quantify the outcome of the

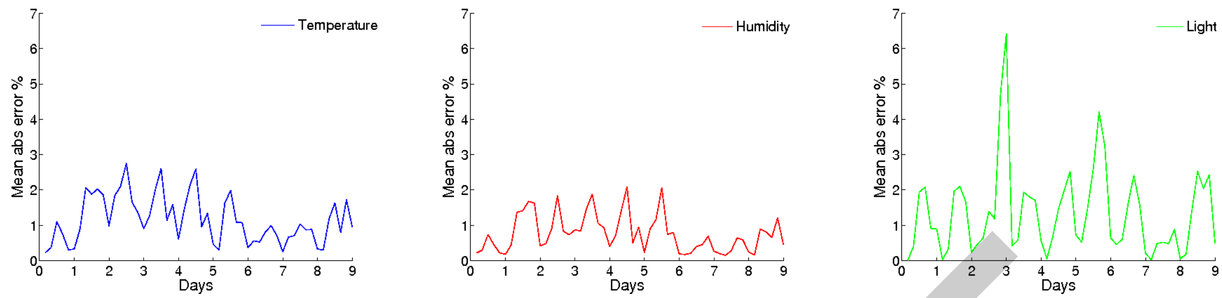

(a) (b) (c)

Figure 5. Mean absolute error for temperature, relative humidity and light exposure in Area 1. In order to compute the average, time slots of 4 hours have been used.

Table I
MEAN ABS ERROR (%)

Area	Temperature	Light	Humidity
1	1.1655	1.3488	0.7961
2	1.2774	1.3560	0.7841
3	0.8318	0.6638	0.5042
4	1.3511	1.6086	0.7658
5	1.1695	1.0098	0.6317

Table II
VARIANCE OF ABS ERROR ($\%^2$)

Area	Temperature	Light	Humidity
1	0.4523	1.5805	0.2814
2	0.4762	1.1046	0.2273
3	0.3420	0.2408	0.1134
4	0.3441	2.2556	0.1784
5	0.3698	0.5470	0.1255

Table III
MAX ABS ERROR (%)

Area	Temperature	Light	Humidity
1	2.7521	6.4209	3.0874
2	2.9180	4.4249	1.8914
3	2.0053	1.9667	1.6625
4	2.5279	6.4697	1.6638
5	2.4300	2.7350	1.4039

proposed algorithm in terms of energy saving. In particular, we estimate the reduction of the overall required sensing and transmissions across all nodes with respect to the original set of raw data.

In order to avoid large prediction errors, we set an upper bound for the sampling rate equal to 8 minutes, whereas the basic sampling rate is slightly over 30 seconds; this results into a sampling frequency ranging from a minimum value of $2,1 \cdot 10^{-3} Hz$ up to a maximum of $33,3 \cdot 10^{-3} Hz$. Figure 6 plots the mean sampling frequency for area 1, for all considered quantities, and Tables IV and V show mean and variance of sampling frequency for each area, respectively. The largest values for the sampling rate (i.e. $4.31 \cdot 10^{-3} Hz$) occur for light, due to the intrinsic nature of such physical quantity, which exhibits lower autocorrelation than temperature and humidity. Nonetheless, such value represents a significant improvement as compared to the basic sampling rate of $33,3 \cdot 10^{-3} Hz$, meaning that, on average, only $\sim$12% samples are needed with respect to the maximum frequency. Temperature and humidity sampling rates are quite close to the minimum sampling rate of $2.1 \cdot 10^{-3} Hz$, thanks to the higher autocorrelation characteristic of these quantities.

In order to quantify the impact of the proposed approach in terms of energy saving, we compared the number of necessary transmissions with respect to the basic approach where each node plainly forwards its samples to the collecting station at the basic, fixed sampling rate.

Table VI shows the percentage of measurements actually needed by our algorithm in order to produce models that are deemed reliable, with respect to the thresholds defined for the prediction error.

We note that the number of samples needed in the worst case amounts to 8.5% for temperature, 6.8% for relative humidity and 12% for light, which is the environmental quantities with higher variance. Energy saving on transmission and sensing is thus greater than 88% in all cases.

V. CONCLUSION AND FUTURE WORKS

This work presented a novel approach aimed to implement energy saving in Wireless Sensor Networks. The use of prediction models for physical quantities allowed to reduce energy consumption of sensor nodes, by adaptively tuning their sampling rate. Experiments carried out on real data showed encouraging results; in particular with a prediction error threshold set to 5%, an overall energy saving amounting to $\sim$90% was obtained. Furthermore our approach allowed us to use a much lower sampling rate than what chosen for the original configuration, while still achieving fully acceptable performances in terms of reliability.

On-going work is being carried out in order to further improve the proposed method. In particular, the behavior of the method needs some refinement in presence of specific data patterns, namely when low-frequency data is followed by high-frequency data. Currently, a high-frequency burst is detected by the base station upon receipt of data from its children; however the transmission of sensor nodes may be scheduled too late, so that the high frequency burst is missed. Such misbehavior can be addressed by computing some kind of frequency measure locally on a sensor node, which is thus allowed to individually decide whether to transmit data before the scheduled time, in order to warn the base station about the fast changes of the environment. Such solution would cause an over-sensing, and might be of interest only with low-power sensors.

Further experiments are also being devised in order to provide an accurate comparison between our method and

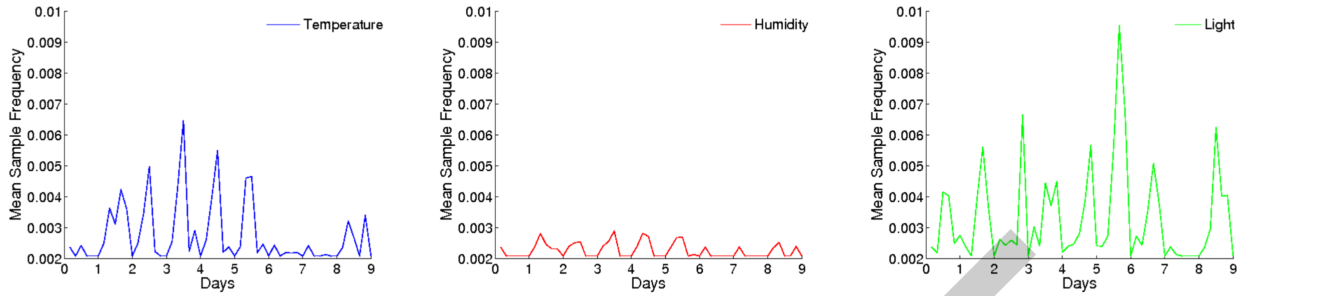

(a) (b) (c)

Figure 6. Sampling rate for temperature, humidity and light exposure in Area 1. In order to compute the average, time slots of 4 hours have been used.

Table IV
MEAN SAMPLING RATE ($Hz \cdot 10^{-3}$)

Area	Temperature	Light	Humidity
1	2.8120	3.3790	2.3340
2	2.9020	3.5690	2.3450
3	2.4870	2.4760	2.2910
4	2.8650	4.3120	2.2990
5	2.7100	3.1110	2.2660

Table V
VARIANCE OF SAMPLING RATE ($Hz^2 \cdot 10^{-6}$)

Area	Temperature	Light	Humidity
1	1.0220	2.4120	0.0580
2	1.2170	2.8990	0.0570
3	0.1990	0.2050	0.0210
4	0.5040	9.1260	0.0380
5	0.3900	1.6400	0.0280

Table VI
FRACTION OF USED SAMPLES (%)

Area	Temperature	Light	Humidity
1	8.1602	9.8903	6.7943
2	8.5154	10.2800	6.7686
3	7.2849	7.2331	6.5291
4	8.2898	12.7218	6.6528
5	7.9044	9.2566	6.6027

non naïve solutions presented in literature; we will make additional experiments taking into account non-fixed accuracy on predictors, in order to infer the connection between energy consumption and model accuracy, and to evaluate the strength of our work with respect to the state of the art.

REFERENCES

[1] G. Anastasi, I. Giannetti, and A. Passarella, "A BitTorrent proxy for Green Internet file sharing: Design and experimental evaluation," *Computer Communications*, vol. 33, no. 7, pp. 794–802, 2010.

[2] D. Estrin, L. Girod, G. Pottie, and M. Srivastava, "Instrumenting the world with wireless sensor networks," in *Proc. of Int. Conference on Acoustics, Speech, and Signal Processing (ICASSP 2001)*, Salt Lake City, Utah, May 2001.

[3] V. Singhvi, A. Krause, C. Guestrin, J. Garrett Jr, and H. Matthews, "Intelligent light control using sensor networks," in *Proc. of the 3rd Intl. Conf. on Embedded networked sensor systems (SenSys)*, 2005, pp. 218–229.

[4] K. Menzel, D. Pesch, B. O'Flynn, M. Keane, and C. O'Mathuna, "Towards a wireless sensor platform for energy efficient building operation," *Tsinghua Science & Technology*, vol. 13, pp. 381–386, 2008.

[5] A. Kansal and M. Srivastava, "An environmental energy harvesting framework for sensor networks," in *Proc. of Symp. on Low Power Electronics and Design.*, 2003, pp. 481–486.

[6] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "Wireless sensor networks: a survey," *Computer networks*, vol. 38, no. 4, pp. 393–422, 2002.

[7] O. Landsiedel, K. Wehrle, and S. Götz, "Accurate prediction of power consumption in sensor networks," in *Proc. of the 2nd IEEE Workshop on Embedded Networked Sensors.*, 2005.

[8] W. Ye, J. Heidemann, and D. Estrin, "An energy-efficient MAC protocol for wireless sensor networks," in *Proc. of IEEE INFOCOM 2002*, vol. 3. IEEE, 2002, pp. 1567–1576.

[9] E. Blass, J. Horneber, and M. Zitterbart, "Analyzing Data Prediction in Wireless Sensor Networks," in *Proc. of IEEE Vehicular Technology Conference.* IEEE, 2008, pp. 86–87.

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

[11] S. Goel, A. Passarella, and T. Imielinski, "Using buddies to live longer in a boring world," in *Proc. of Pervasive Computing and Communications (PERCOM) Workshops, 2006.* IEEE, 2006, pp. 5–346.

[12] D. Tulone and S. Madden, "PAQ: Time series forecasting for approximate query answering in sensor networks," *Wireless Sensor Networks*, pp. 21–37, 2006.

[13] M. Mollanoori, M. Hormati, and N. Charkari, "An Online Prediction Framework for Sensor Networks," in *16th Iranian Conference on Electrical Engineering*, 2008.

[14] A. Dempster, N. Laird, D. Rubin *et al.*, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 39, no. 1, pp. 1–38, 1977.

[15] S. Madden, "Intel Lab Data," http://db.csail.mit.edu/labdata/labdata.html.

[16] A. De Paola, G. Lo Re, F. Milazzo, and M. Ortolani, "Adaptable Data Models for Scalable Ambient Intelligence Scenarios," in *Proc. of the 25th Intl Conf. on Information Networking*, 2011.

[17] C. Guestrin, P. Bodik, R. Thibaux, M. Paskin, and S. Madden, "Distributed regression: an efficient framework for modeling sensor network data," in *Proc. of ACM IPSN'04.* ACM, 2004, pp. 1–10.