

Online Energy Harvesting Prediction in Environmentally-Powered Wireless Sensor Networks

Alessandro Cammarano, Chiara Petrioli, *Senior Member, IEEE*, and Dora Spenza, *Member, IEEE*

Abstract—The increasing popularity of micro-scale power-scavenging techniques for Wireless Sensor Networks (WSNs) is paving the way to energy-autonomous sensing systems. To sustain perpetual operations, however, environmentally-powered devices must adapt their workload to the stochastic nature of ambient sources. Energy prediction models, which estimate the future expected energy intake, are effective tools to support the development of proactive power management strategies. In this work, we present Pro-Energy, an energy prediction model for multi-source Energy-Harvesting WSNs that leverages past energy observations to forecast future energy availability. We then propose Pro-Energy-VLT, an extension of Pro-Energy that combines our energy predictor with timeslots of variable lengths to adapt to the dynamics of the power source. To assess the performance of our proposed solutions, we use real-life solar and wind traces, as well as publicly-available traces of solar irradiance and wind speed. A comparative performance evaluation shows that Pro-Energy significantly outperforms state-of-the-art energy predictors, by improving the prediction accuracy of up to 67%. Moreover, by adapting the granularity of the prediction timeslots to the dynamics of the energy source, Pro-Energy-VLT further improves the prediction accuracy, while reducing the memory footprint and the energy overhead of energy forecasting.

I. INTRODUCTION

Energy harvesting allows to use energy from the environment to power embedded devices and nodes of Wireless Sensor Networks (WSNs) [1], [2]. By scavenging energy from their surroundings, energy-harvesting wireless sensor nodes can significantly increase their typical lifetime: If the harvested energy is efficiently utilized, low-power devices can last virtually forever. However, although potentially unlimited, the energy provided by ambient power sources is neither constant nor always available. For instance, solar-powered nodes experience significant changes in the power harvested over time, due to the diurnal cycle in solar energy, varying weather conditions and seasonal patterns. Moreover, the position of the nodes and the orientation of their solar cell strongly impact on their energy intake: even if two nodes are physically co-located, their harvesting rates may significantly differ [3]. The uncertainty in the energy availability provided by ambient sources raises new challenges in developing reliable and energy-efficient power-management solutions for Energy-Harvesting WSNs (EH-WSNs) [1]. For this reason, many works targeting

environmentally-powered systems require the energy harvesting profile to be known in advance [4], [5] or that accurate predictions of the future energy intake are available [6]. In the case of predictable energy sources, such as solar light, energy prediction models can estimate the expected energy intake in the near future. Forecasting future power income allows to exploit the available energy at best, minimizing both periods in which nodes are not active due to lack of energy, and waste of energy in periods during which the harvesting rate is high. An increasing number of works on EH-WSNs thus rely on the use of online energy prediction methods to implement harvesting-aware solutions, such as algorithms for smart energy allocation and spending [7], communication protocols [8] and power management strategies based on dynamic load adaptation [9], [10].

In this paper we present a framework for accurate online prediction of the energy intake for multi-source (i.e., solar and wind) energy-harvesting wireless sensor nodes. In particular, we make the following contributions:

- We present Pro-Energy (PROfile Energy prediction model), an energy prediction model that leverages past energy observations for accurate estimations of future energy availability at both short (few minutes to half an hour) and medium (a few hours) predictions horizons.
- We propose Pro-Energy-VLT (PROfile Energy prediction model with Variable-Length Timeslots), which combines Pro-Energy with timeslots of variable lengths, whose granularity is set coarser or finer based on the dynamics of the power source. Thanks to the online timeslots adaptation, Pro-Energy-VLT better captures patterns of the harvesting process, improving accuracy and reducing the memory and energy overhead of energy prediction.
- We evaluate the performance of our proposed predictors in several scenarios, using both real-life harvesting traces that we obtained by interfacing Telos B nodes with photovoltaic panels and wind micro-turbines, and publicly-available datasets of solar and wind energy availability.
- We present a case study for energy prediction in WSNs and assess the practical feasibility of our proposed approach through implementation on solar-powered motes.

The remainder of this paper is organized as follows. Related work is surveyed in Section II. We present our prediction model, Pro-Energy, in Section III. In Section IV a variant of Pro-Energy using timeslots of variable length, named Pro-Energy-VLT, is proposed. A comparative performance evaluation of Pro-Energy and of Pro-Energy-VLT against state-of-the-art energy prediction models is performed in Section V. Section VI concludes the paper.

A. Cammarano, C. Petrioli and D. Spenza are with the Department of Computer Science, University of Rome “La Sapienza”.
E-mail: {cammarano, petrioli, spenza}@di.uniroma1.it.

A portion of this paper is an extension of [23], [24]. The research contribution presented in this paper has been supported by the projects *Smartour (Intelligent platform for tourism)* and *Social Museum & Smart Tourism* founded by MIUR, and by the project *Harvesting-Aware Communication Protocols in the Self-Powered Internet of Things* founded by University of Rome “La Sapienza”. The authors also thank the MICREL lab, University of Bologna, for sharing their wind dataset.

II. RELATED WORK

Driven by the rapid diffusion of small-scale energy-harvesting techniques for embedded devices, and by the challenges posed by their uncertain power availability, many recent works have been focusing on energy prediction techniques for energy-harvesting wireless sensor networks (EH-WSNs). Kansal et al. were the first to present a solar energy prediction scheme for EH-WSNs: the Exponentially Weighted Moving-Average (EWMA) [11]. EWMA, based on an exponentially weighted moving-average filter, relies on the assumption that the energy available at a given time of the day is similar to the energy harvested at the same time on the previous days. The amount of energy available during the past days is maintained in EWMA as a weighted average, in which the contribution of older data is exponentially decaying. EWMA provides quite accurate results in presence of scarce weather variability, but its performance significantly degrades when weather conditions are frequently changing (e.g., when a mix of sunny and cloudy days occurs in a row). To address this shortcoming, Piorno et al. proposed WCMA, a Weather-Conditioned Moving Average prediction model [12], [13]. Similarly to EWMA, WCMA takes into account the energy harvested in the previous days, but it additionally introduces a weighting factor to quantify how the weather conditions of the current day changed with respect to the previous days. In case of frequently changing weather conditions, WCMA is shown to obtain average prediction errors almost 20% smaller than EWMA. The solution proposed by Noh et al. in [14] uses the EWMA model to keep track of the solar energy profile observed in the past. With respect to EWMA, an additional scaling factor is introduced to adjust future energy expectations based on short-term varying weather conditions. Scaling is performed by computing the ratio between the actual energy harvested during the current timeslot and the energy predicted for the same timeslot. Linear interpolation is then used to reduce the scaling value for slots that are further ahead in time. In [15], Lu et al. addresses energy-harvesting prediction for real-time embedded systems by investigating common techniques for real time series prediction. Regression analysis is shown to obtain the best accuracy for energy predictions within a time horizon of 1 second, but performance on short and medium term prediction horizons are not discussed.

These energy prediction methods base their forecasting upon locally-collected data. In application scenarios in which Internet connectivity is available, e.g., by means of an external server connected to the WSN, energy prediction models that can take advantage of external global information have also been proposed. For example, Renner investigated in [16] the integration of local information with global cloud cover forecasting to improve long-term predictions (of at least one day) in solar-powered systems. A system architecture that obtains global weather forecasts and disseminates them into the sensor network is presented in [17]. To reduce the additional complexity and overhead introduced by the need to periodically disseminate weather forecasting updates in the network, Renner and Nguyen proposed a lossless compression method for cloud-cover forecasts [18]. Sharma et al.

presented in [19] an energy-harvesting prediction method for solar and wind powered sensor systems, which uses models of solar panels and wind turbines to convert global weather forecasting, i.e., cloud cover and wind speed, into energy harvesting predictions. Their approach is shown to outperform energy predictors based only on past observations for long-term prediction horizons between 3 hour to 3 days.

Another class of solutions relies on machine learning techniques for energy-harvesting predictions. For example, Lu and Whitehouse presented in [20] SunCast, a system for using natural sunlight inside a building to reduce the electricity demand of artificial lighting. SunCast consists of a control algorithm to provide stable indoor lighting levels and of a sunlight prediction algorithm to generate a distribution of predicted sunlight values. Being not devised to be locally run on wireless sensor nodes, the computational and memory demands of SunCast may exceed the limited resources typically available on WSN nodes. Sharma et al. evaluated in [21] multiple regression techniques to automatically create site-specific prediction models for solar power generation from weather forecasts. Such predictions models are intended to be used by a smart grid and by individual smart houses to plan electricity generation and consumption in advance.

As for the use of variable-length timeslots, Renner and Turau were the first to introduce an adaptive slotting scheme for energy forecasting in EH-WSNs [22]. Their proposed approach improves the forecast accuracy by identifying the actual pattern of energy intake at the end of each day. An online algorithm is presented to derive the optimal timeslot distribution, defined as the daily timeslot division that minimizes the representation error between the harvester output at time t and the mean of the harvesting samples of the timeslot covering time t . Due to the difficulty of computing the optimal timeslot distribution on resource-limited WSN nodes, an adaptive slot distribution algorithm is also discussed.

III. PRO-ENERGY

In this section, we present Pro-Energy [23], [24], an energy prediction algorithm for wireless sensor networks that uses past-days observations to derive predictions on the future energy intake. The key concept of Pro-Energy is to make use of harvested profiles representing the energy intake available during different types of “typical” days. For instance, days may be classified into sunny, cloudy or rainy, and a characteristic profile may be associated to each of these categories. Pro-Energy works as follows. Each day is discretized into a number N of non-overlapping timeslots of equal length. Predictions are performed once per slot. The energy received during the current day is stored in a vector, C , of length N , containing the energy obtained during each of the past timeslots. A pool of energy profiles observed in the past is also maintained in a matrix E , of size $D \times N$. These profiles represent the energy obtained during a number D of typical days, which are used by Pro-Energy to forecast future energy intake over short and medium term time frames. Pro-Energy is made up of three logical components:

TABLE I
TABLE OF NOTATIONS

Symbol	Explanation
N	Number of timeslots in a day
D	Number of past energy profiles stored
E	Matrix of stored energy profiles ($D \times N$)
E_t^s	Harvested energy observed during timeslot t of stored day s
C	Energy harvested during the current day
C_t	Energy harvested during timeslot t of the current day
\hat{E}_t	Predicted energy at timeslot t on the current day
K	Number of past observations used to compute profiles similarity
α	Weighting factor for short-term prediction
F	Prediction horizon (number of time slots)
G	Correlation parameter
P	Number of energy profiles combined for energy predictions
WP	Weighted profile (combination of P profiles)
γ	Weighting factor for short and medium-term predictions
MAE	Mean Absolute Error
$MADP$	Mean Absolute Deviation Percent
p_t	Average harvesting power during timeslot t
d_t	Duration of timeslot t (in seconds)
$E_{\mathcal{H}}$	Energy harvested during the time interval $[\tau, \tau + \mathcal{H}]$
$\hat{e}_{\mathcal{H}}$	Energy prediction for the time interval $[\tau, \tau + \mathcal{H}]$

- 1) the **profile analyzer**, which selects the stored profile that is the most similar to the current day among those in the pool (run at the end of each slot, N times per day);
- 2) the **prediction module**, which computes future predictions for both short and medium term prediction horizons (run at the end of each slot, N times per day);
- 3) the **profile pool refresher**, which periodically updates the pool of stored energy profiles by discarding profiles that have become obsolete (run at the end of each day, once per day).

The profile analyzer and the prediction module are invoked at the beginning of each time slot to derive future energy predictions, while the profile pool refresher is run at the end of each day to update the pool of stored energy profiles.

A. Profile analyzer

Pro-Energy computes energy predictions by combining the information about the energy harvested during the current day with the energy intake obtained during the stored profile that is the most similar to the current day. For each stored profiles, similarity with the current day is calculated as the mean absolute error (MAE) of the energy harvested over the last K timeslots. More formally, we indicate as E_t^d the stored profile that is the most similar to the current day C among the D typical profiles stored in the pool. At each timeslot t , E_t^d is computed as follows:

$$E_t^d = \min_{E_i^s \in E} \sum_{i=t-K+1}^t \frac{1}{K} |C_i - E_i^s| \quad (1)$$

Similarity is computed over the last K timeslots. In addition to reducing the computational overhead, using only the last K timeslots (rather than the entire profile) allows to react more quickly to situations in which the weather conditions change during the current day. For instance, if during the current day

a sunny morning is followed by a cloudy afternoon, the profile analyzer will likely select different profiles in the morning and in the afternoon based on the current weather conditions.

B. Prediction module

Pro-Energy can be used to provide future energy predictions over both short (few minutes to half an hour) and medium (a few hours) prediction horizons. In the following, we first detail the approach used in Pro-Energy for short-term energy predictions, and then explain how the short-term predictor is generalized for prediction horizons of up to a few hours.

1) *Short-term energy predictions*: When computing energy predictions, Pro-Energy tries to match the observations of the current day with one of the typical profiles stored in its pool. Specifically, Pro-Energy computes the predicted value for the next timeslot as a linear combination of the energy observed in the last timeslot C_t and the value for the next timeslot reported in the stored profile. More formally, being E_t^d the stored profile that at timeslot t is the most similar to the current day, the predicted energy intake for the next slot, $t+1$, of the current day is computed as:

$$\hat{E}_{t+1} = \alpha \cdot C_t + (1 - \alpha) \cdot E_{t+1}^d \quad (2)$$

where:

- \hat{E}_{t+1} is the predicted energy intake in timeslot $t+1$ of the current day;
- E_{t+1}^d is the energy harvested during timeslot $t+1$ on the stored day d ;
- C_t is the energy harvested during timeslot t on the current day C ;
- α is a weighting factor, $0 \leq \alpha \leq 1$.

Parameter α weights the combination between the value reported in the stored profile and the current energy observation, i.e., the energy observed in the last slot, C_t .

2) *Medium-term energy predictions*: When computing short-term predictions, considering the correlation between two consecutive timeslots usually helps increasing the prediction accuracy. This approach, however, may not be as effective in case of medium-term predictions. In fact, the correlation between the energy observed at time t and the one observed at time $t+\delta t$ generally decreases for increasing δt . To quantify such correlation, we analyzed both solar and wind harvested data, discretizing days into 48 timeslots of 30 minutes each, and computing the correlation between the average energy harvested during timeslot t and that obtained during timeslot $t+\delta t$, $\forall t = \{0, \dots, 48\}$, $\forall \delta = \{0, \dots, 48\}$. Figure 1 shows an example of such an analysis, reporting the Pearson correlation¹ computed for $t = 8:30$ AM (results are similar for different timeslots). In case of solar energy harvesting, there is a *strong* correlation (i.e., correlation coefficient > 0.7) between the energy harvested at 8:30 AM and that gathered during the successive 4-6 slots (2-3 hours) (Fig. 1(a)). Wind is generally more variable than solar: energy harvested at 8:30 in

¹The Pearson correlation coefficient ranges from -1 to 1. An absolute value of 1 implies a linear relationship between the energy harvested at two different timeslots, while a value of 0 means there is no linear correlation between them.

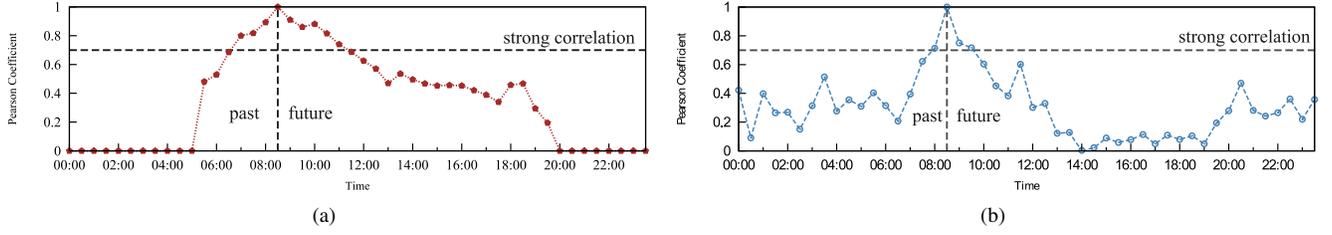


Fig. 1. Pearson autocorrelation coefficient for different energy sources: (a) solar and (b) wind.

the morning shows *strong* correlation only with the successive 1-2 slots, and weak correlation afterwards (Fig. 1(b)).

Knowledge of correlation can be exploited for better prediction accuracy. To this purpose, we introduce a new parameter, γ , which determines the influence of the last energy observation while computing predictions for the next F future slots. Assuming that two slots at a distance equal to or greater than G show only a weak correlation, the γ parameter to be used when predicting the energy intake for the future slot at distance i from the current slot is defined as:

$$\gamma_i = \begin{cases} \alpha \cdot \left(1 - \frac{i-1}{G}\right), & \text{if } i \leq G \\ 0 & \text{if } i > G \end{cases} \quad \forall i, 1 \leq i \leq F \quad (3)$$

where:

- α is the weighting factor defined in Equation (2);
- i is the i^{th} timeslot in the future, with respect to the current slot, t ;
- G is the number of timeslots in the future whose correlation with timeslot t is above a given threshold;
- F is the number of future timeslots for which Pro-Energy is computing energy predictions.

The γ_i parameter plays a similar role in Equation (4) (defined in the following) as the weighting parameter α in Equation (2): It allows to combine the energy value of the stored profile with the current energy observation. However, the weight associated to the value observed during the current slot progressively decreases when computing predictions for timeslots that are further away in time. In particular, such a weight is set to zero for timeslots that are more than G slots far in the future, as there is weak or no correlation between the energy harvested at timeslot t and at timeslot $t + G$.

Having defined such γ_i parameter, the medium term predictions are then computed by using a generalization of the short-term version (Equation 2):

$$\hat{E}_{t+i} = \gamma_i \cdot C_t + (1 - \gamma_i) \cdot E_{t+i}^d \quad (4)$$

Short-term predictions are thus computed according to Equation 4, where the prediction horizon F is equal to 1, $i = 1$ and $\gamma_i = \alpha$.

C. Stored profiles updates

Pro-Energy maintains a pool of D typical profiles, each ideally representative of a different weather condition. In order to adapt predictions to changing seasonal patterns, this pool has to be periodically updated. In particular, at the end of each day the pool of stored profiles can be updated with the profile observed during the current day, C . Pro-Energy jointly implements two update strategies:

- 1) If there is a profile in the pool that was stored longer than x days ago, substitute it with the profile observed during the current day, C .
- 2) If two profiles in the pool, E^{d_1} and E^{d_2} , are very similar, i.e., their MAE is below a given threshold T_s :

$$\frac{1}{N} \sum_{i=1}^N \left| E_i^{d_1} - E_i^{d_2} \right| < T_s \quad (5)$$

substitute with C the one among the two that is the most similar to the current day. In case of multiple pairs of similar profiles, the most similar to C is selected.

The first strategy allows to discard profiles that have become obsolete, while the second strategy allows to maintain a pool of profiles that are ideally representative of different weather conditions, by discarding profiles that are very similar.

D. Combination of multiple profiles

Pro-Energy predictions can be further improved through a technique that allows to combine multiple profiles together. Such method selects a set of P profiles, instead of a single one, among the D profiles stored in the E matrix and combines them to form a “weighted” profile WP .

The rationale behind the use of multiple profiles is to consider different possible evolutions of the current day. For instance, a sunny morning may be followed by a cloudy or rainy afternoon. While computing medium and long term predictions, considering a single profile may lead to poor accuracy if significant variations occur in the weather conditions. On the contrary, considering multiple profiles allows to account for these potential changes, reducing the prediction error at the price of a small additional overhead.

Let $E^{d_1}, E^{d_2}, \dots, E^{d_P}$ be the ordered list of profiles that are most similar to the current day C , i.e., profiles with the smaller Mean Absolute Errors. The weighted profile WP , for the future slot $t + i$, $i \in \{1, 2, \dots, F\}$, is computed as:

$$WP_{t+i} = \frac{1}{P-1} \sum_{j=1}^P w_j \cdot E_{t+i}^{d_j} \quad (6)$$

where

$$w_j = 1 - \frac{MAE_k(E^{d_j}, C)}{\sum_{j=1}^P MAE_k(E^{d_j}, C)} \quad (7)$$

Energy predictions for the future slot $t+i$ are then computed based on both the value for such slot stored in the WP profile and the energy harvested during the last time slot:

$$\hat{E}_{t+i} = \gamma_i \cdot C_t + (1 - \gamma_i) \cdot WP_{t+i} \quad (8)$$

where:

\hat{E}_{t+i} is the predicted energy in timeslot $t+i$ for the current day;

C_t is the harvested energy during last timeslot;

WP is the vector of the combination of timeslot $t+i$ of the P profiles;

γ_i is a correlation factor for prediction of slot $t+i$.

Equation (8) is a generalization of Equation (4). Specifically Equation (4) is obtained by Equation (8) by using only one profile for future energy predictions.

IV. PRO-ENERGY WITH VARIABLE-LENGTH TIMESLOTS

Similarly to other state-of-the-art predictors, Pro-Energy divides each day into N equal-length timeslots to reduce the time and memory overhead of predictions. Although very common, however, using equally-distributed timeslots allows to update future energy estimations only at pre-defined instants in time. This is generally an under-performing strategy, as the variability of the harvesting source is typically not constant over time. A real-life example is shown in Fig. 2, which reports an energy harvesting trace we collected by interfacing a Telos B mote with a XOB17-04x3 solar cell [25]. The dynamics of the energy source can be better captured by using timeslots with different granularity, so that energy predictions are updated more or less frequently based on how rapidly the current energy-harvesting rate is changing. Based on this observation, we propose Pro-Energy-VLT (PROfile Energy prediction model with Variable-Length Timeslots), an energy prediction model that combines Pro-Energy with timeslots of variable lengths. Pro-Energy-VLT dynamically resizes the prediction timeslots so as to obtain a timeslots granularity that is coarser (during periods of low dynamicity) or finer (during periods of high dynamicity) based on the dynamics of the power source. An online algorithm is locally run by each node to periodically update both the number of timeslots and their size. Pro-Energy-VLT works as follows. During the initial setup phase, each day is divided into N equal-length timeslots, as in Pro-Energy. A weight is then assigned to each timeslot t , based on a simple yet effective heuristic that estimates how rapidly the availability of the energy source changes over different timeslots (Section IV-A). Higher weights are assigned to timeslots that cover periods of time during which the availability of the energy source varies sharply, while lower weights are associated with timeslots that cover periods of time during which the energy source provides stable levels of power. Then, Pro-Energy-VLT redistributes the N timeslots based on their weights (Section IV-B). As a first step, the contiguous timeslots having a weight equal to zero (or, more generally, below a minimum threshold) are identified and merged together. For example, in the solar energy harvesting case, a single large timeslot is created to cover the whole night, by merging together z timeslots. The remaining $N-1$ timeslots are then distributed proportionally to their weights. In this process, timeslots with higher weights are divided into a higher number of sub-timeslots, thus resulting in a finer timeslot granularity. Timeslots with low weights are instead merged with neighbor timeslots, leading to a coarser

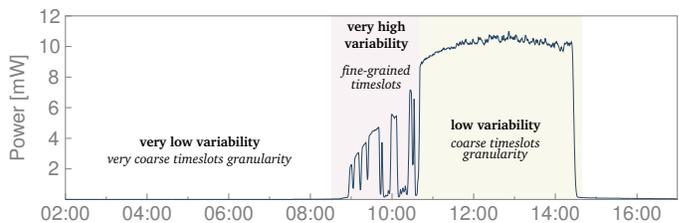


Fig. 2. Energy harvesting trace collected by interfacing a Telos B mote with a XOB17-04x3 solar cell.

timeslots granularity. As a last step, the pool of stored profiles maintained by Pro-Energy-VLT is updated to reflect the new timeslot setting. To this end, the energy harvesting profiles stored in the pool are interpolated to determine the power values corresponding to the new timeslots setting (Section IV-C).

This timeslot adapting phase is periodically run by each node. Every D days, Pro-Energy-VLT uses the information collected over the past D days to produce a new timeslot setting, in which the N timeslots used for prediction are re-distributed based on the dynamics of the energy source.

A. Weights assignment

During the weight assignment phase, each timeslot t is assigned a weight based on the variability of the harvesting process during timeslot t with respect to the previous timeslot $t-1$. To quantify such a variation, the difference between the average harvested power during timeslots t and $t-1$ is considered. Since in Pro-Energy-VLT timeslots generally have different durations, weights assignment also takes into account the size of the timeslot, d_t . A weight is assigned to each timeslot t according to the following heuristic:

$$w_t = \log(d_t \text{abs}(p_t - p_{t-1}) + 1) \quad (9)$$

where p_t and p_{t-1} denote the average power harvested during timeslots t and $t-1$, respectively, and d_t is the duration of timeslot t in seconds. (Table I). The logarithm function is used to reduce the absolute value of the assigned weights, while preserving their relative difference.

B. Timeslot resizing

During the initial setup phase, Pro-Energy-VLT merges together one or multiple sets of contiguous timeslots having a weight equal to zero (or below a minimum threshold). Let z be the number of generated timeslots during this initialization phase, the remaining $N-z$ timeslots are then distributed proportionally to their weights. In particular, each timeslot t is split into a number n_t of sub-timeslots, as follows:

$$n_t = \left\lfloor \frac{w_t}{\sum_t w_t} \cdot (N-z) \right\rfloor, \quad (10)$$

where w_t is the weight assigned to t during the weight assignment phase described in Section IV-A. If the value of n_t is equal to zero (or lower than a small threshold), timeslot t covers a period of time during which the energy source provides stable levels of power. A coarser timeslots granularity is thus used during this period by merging timeslot t with the successive timeslots.

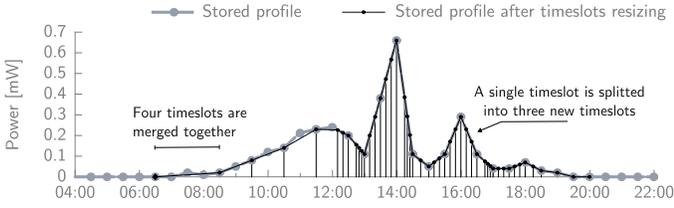


Fig. 3. Example: an energy profile stored in the Pro-Energy-VLT pool is updated to new timeslots settings.

C. Profile pool update

As a last step, the pool of energy profiles maintained by Pro-Energy-VLT must be updated to reflect the new timeslot setting. To this end, the energy harvesting profiles stored in the pool are linearly interpolated to determine the power values corresponding to the new timeslots setting.

Figure 3 shows an example of a profile stored in the Pro-Energy-VLT pool that is updated after timeslots resizing.

V. PERFORMANCE EVALUATION

In this section, we systematically evaluate the performance of both Pro-Energy and Pro-Energy-VLT in different settings, using six datasets of traces of harvested energy:

- **solar-ROME, wind-ROME**: real-life solar and wind traces obtained from a testbed in Rome, Italy;
- **wind-BO**: real-life wind dataset obtained from a testbed in Bologna, Italy;
- **indoor-EnHANTs**: database of indoor radiant light measurements from the EnHANTs (Energy Harvesting Active Networked Tags) project [26];
- **solar-ORNL, wind-ORNL**: traces of solar and wind availability obtained from the National Renewable Energy Laboratory (NREL) at Oak Ridge, Tennessee [27].

We obtained real-life solar and wind traces (datasets solar-ROME and wind-ROME) by interfacing Telos B motes with photovoltaic cells (Fig. 4(a)) and with wind micro turbines (Fig. 4(b)). The Telos B is a low-power wireless sensor node featuring a typical current consumption is of just $5.1\mu\text{A}$ in sleep mode with the MCU in standby and the radio off. In active mode (MCU on and radio off) its typical current consumption is of 1.8mA . When the radio is in RX state, the Telos B consumes around 21.8mA . The energy harvesting board shown by Fig. 4(a) stores energy harvested from the solar panel using rechargeable batteries or supercapacitors. In our experiments, we used a POW112D2P solar panels of size $7 \times 5.5 \text{ cm}$ providing a typical peak power of 450 mW , 50F Panasonic Gold supercapacitors and Duracell NiMH 2450mAh rechargeable batteries. The harvesting board includes a maximum power point controller to maximize the harvesting efficiency [28]. The TPS63001 DC/DC converter from Texas Instruments is used to regulate the output voltage. The voltage of the solar panel and of the supercapacitors of the node are sampled using two ADC input ports of the Telos B mote (Fig. 5). A direct voltage look up table is used by the node to estimate the amount of power generated by the cell based on its output voltage. We developed a dedicated TinyOS [29] application to track the amount of energy generated by solar

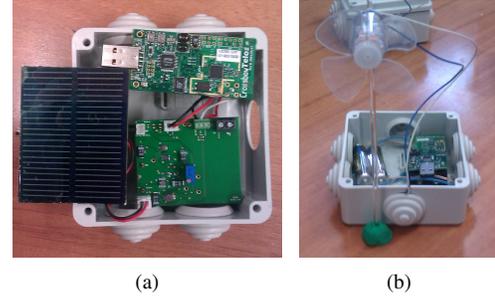


Fig. 4. Telos B motes interfaced with (a) a photovoltaic cell; and (b) a wind micro turbines.

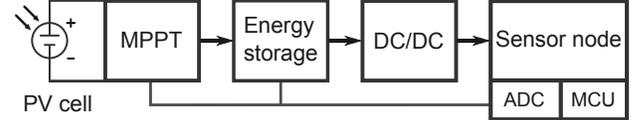


Fig. 5. Simplified block diagram of the solar harvesting board.

and wind harvesters every 30 seconds. The monitoring motes were deployed outdoors, on the roof terrace of University of Rome “La Sapienza” CS Dept. building in downtown Rome. We also obtained additional real-life wind harvesting data from an outdoor testbed in Bologna, Italy, of Telos B motes equipped with wind micro turbines that collected data for 75 days. The fourth type of energy traces we use is a database of indoor radiant light measurements collected in office buildings in New York City within the EnHANTs (Energy Harvesting Active Networked Tags) project of Columbia University [26], consisting of 90 days of measurements. The other two datasets, i.e., solar-ORNL and wind-ORNL, have been obtained from the National Renewable Energy Laboratory (NREL) at Oak Ridge, Tennessee, and consist of 90 days of solar and wind data with a granularity of one sample per minute. To use such traces in our performance evaluation, we converted raw weather data, i.e., irradiance and wind speed values, into energy harvesting estimations. We calculate the power P_s harvested by a solar cell of size A and efficiency η as:

$$P_s = A \cdot \eta \cdot I, \quad (11)$$

where I is the radiant energy incident onto surface. Based on characteristics of real-life micro photovoltaic cells, we set $\eta = 0.15$ and $A = 7 \times 5 \text{ cm}^2$ to obtain the harvested power from indoors EnHANTs traces and $\eta = 0.17$ and $A = 22 \times 7 \text{ mm}^2$ to obtain the harvested power from outdoors NREL traces. As for wind energy harvesting, we estimate the output power P_w of the wind micro turbine as follows [30]:

$$P_w = \frac{1}{2} \cdot v^3 \cdot A \cdot \rho \cdot C_p, \quad (12)$$

where v is the wind speed in m/s , A is the rotor swept area in m^2 , ρ is the air density (typically 1.25 kg/m^3), and C_p is the power coefficient, which represents the ratio of power extracted by the turbine to the total contained in the wind. We consider a wind micro turbine with a rotor diameter of 5 cm , and set the power coefficient C_p to 1.5% in order to obtain harvesting power values consistent with real-life measurements. Simulations are run using GreenCastalia [31],

an open-source extension we have developed for the popular Castalia simulator [32] to model networks of embedded devices with energy-harvesting capabilities. GreenCastalia allows to perform simulations using timestamped power traces collected through real-life deployments, and it includes models of energy harvesting devices to obtain energy harvesting estimations from raw weather data traces.

A. Prediction algorithms evaluation

We evaluated the performance of Pro-Energy by comparing the amount of energy actually harvested during each timeslot against the predicted energy intake. To this end, we introduce a new metric of comparison that evaluates prediction errors from the point of view of a generic external module exploiting future energy intake predictions (e.g., an application-level component or a harvesting-aware communication protocol). In the general case, such a module may require estimates of future energy intake asynchronously with respect to the predictor timeslots. For example, harvesting-aware task allocation algorithms, such as [33], use energy prediction to support task allocation decisions, thus requiring predictions whenever a new task arrives, rather than at pre-defined instants in time. Moreover, the prediction horizon may be equal to the task duration, rather than being a multiple of the prediction timeslot. To handle such requirements, differently from previous works, we estimate the prediction error over time windows that are not necessarily related to the actual size of prediction timeslots. We evaluate the accuracy of different prediction algorithms in a general setting, by considering different time horizons, i.e., 5, 10, 20, 30, 60 and 120 minutes. For generality, we assume the external module requires energy predictions for all the future prediction horizons every 5 minutes. The overall error of a prediction algorithm is thus computed as Mean Absolute Deviation Percent (MADP)² of all delivered predictions over the whole dataset, as follows:

$$MADP = 100 \cdot \frac{\sum |E_{\mathcal{H}} - \hat{e}_{\mathcal{H}}|}{\sum E_{\mathcal{H}}}, \quad (13)$$

where $E_{\mathcal{H}}$ is the actual energy harvested during the time interval $[\tau, \tau + \mathcal{H}]$, and $\hat{e}_{\mathcal{H}}$ is the energy prediction for the same interval, which is computed as follows (Equation 14). Each time interval $[\tau, \tau + \mathcal{H}]$ is covered by a number n of prediction timeslots, e.g., t_j, \dots, t_{j+n-1} . The predicted energy intake for a time horizon \mathcal{H} , i.e., the total energy that is expected to be harvested in the time interval $[\tau, \tau + \mathcal{H}]$, is computed as:

$$\hat{e}_{\mathcal{H}} = \sum_{i=j}^{n-1} d_i \times p_i, \quad (14)$$

where d_j is the duration of each timeslot t_j within the time interval $[\tau, \tau + \mathcal{H}]$, and p_j is the expected average harvested power during timeslot t_j .

²The MADP metric, which is not calculated on a daily basis, correctly takes into account errors introduced by periods of very low energy intake. Other commonly-used metrics, such as the Mean Absolute Percentage Error, require to filter out such periods [23], [34] to prevent the introduction of large errors during timeslot with very low energy intake.

B. Pro-Energy: Evaluation results

We compare the performance of Pro-Energy with that of four state-of-the-art energy predictors: EWMA, the energy prediction model proposed by Noh and Kang in [14], which we denote for brevity AEWMA, WCMA and the sunlight prediction algorithm of SunCast. In our experiments, we set $N = 48$. WCMA, AEWMA and EWMA assume the energy harvesting in the near future to be related to the energy intake at the same time on the previous days. While EWMA and AEWMA maintain historical data as a single vector of size N , WCMA stores a matrix of size $D \times N$, where D is the number of previous days used for energy predictions. SunCast takes into account for predictions all the historical data previously collected. Figure 6 shows the prediction error (computed according to Equation (13)) of the considered prediction algorithms for different solar and wind datasets and for increasing prediction horizons. To perform a fair comparison, we set the coefficients of each prediction model to their optimal value, i.e., the ones minimizing the overall MADP error. Pro-Energy consistently outperforms EWMA, AEWMA, WCMA and SunCast for both short and medium term solar energy predictions. In case of solar energy predictions with a prediction horizon of up to 2 hours, Pro-Energy performs up to 49% (Fig. 6(b)) and up to 67% (Fig. 6(c)) better than competing schemes. Pro-Energy also achieves a good accuracy for indoor light forecasting. For the indoor-EnHANTs dataset (Fig. 6(d)), using Pro-Energy results in a 32%-48%, 9%-15%, 1%-17% and 50%-60% lower MADP errors than EWMA, AEWMA, WCMA and SunCast, respectively. Results are similar for wind harvesting forecasting. In case of wind energy predictions with a prediction horizon of up to 2 hours, Pro-Energy performs up to 35% (Fig. 6(e) and Fig. 6(c)) and up to 53% (Fig. 6(f)) better than competing schemes, depending on the considered wind dataset. Overall, WCMA shows a good accuracy in case of short-term energy prediction, but its performance tends to degrade in case of medium-term predictions. This is due to the fact that WCMA, being designed to only deliver predictions for the next timeslot, does not exploit the correlation between current observations and future ones to adjust energy predictions over medium-term forecasting horizons. Performance of AEWMA significantly vary for different datasets. This is because AEWMA applies to future energy prediction a scaling factor that is based on the ratio between the energy predicted for a timeslot and the actual energy harvested during that timeslot. Depending on the considered energy source and dataset, this scaling factor may be applied to timeslots that have very weak or no correlation with the current timeslots, which reduces prediction accuracy. SunCast shows good accuracy in case of solar light predictions (Fig. 6(c)), but it suffers from high MADP errors when predicting wind availability and indoor irradiance levels. The prediction accuracy of EWMA is generally limited, as it does not use information about the current energy intake to adjust prediction for future slots: at the end of each slot t , the only prediction updated is that of slot t , which will be used for energy prediction on the next day. This leads to EWMA reacting slowly to changes in weather conditions.

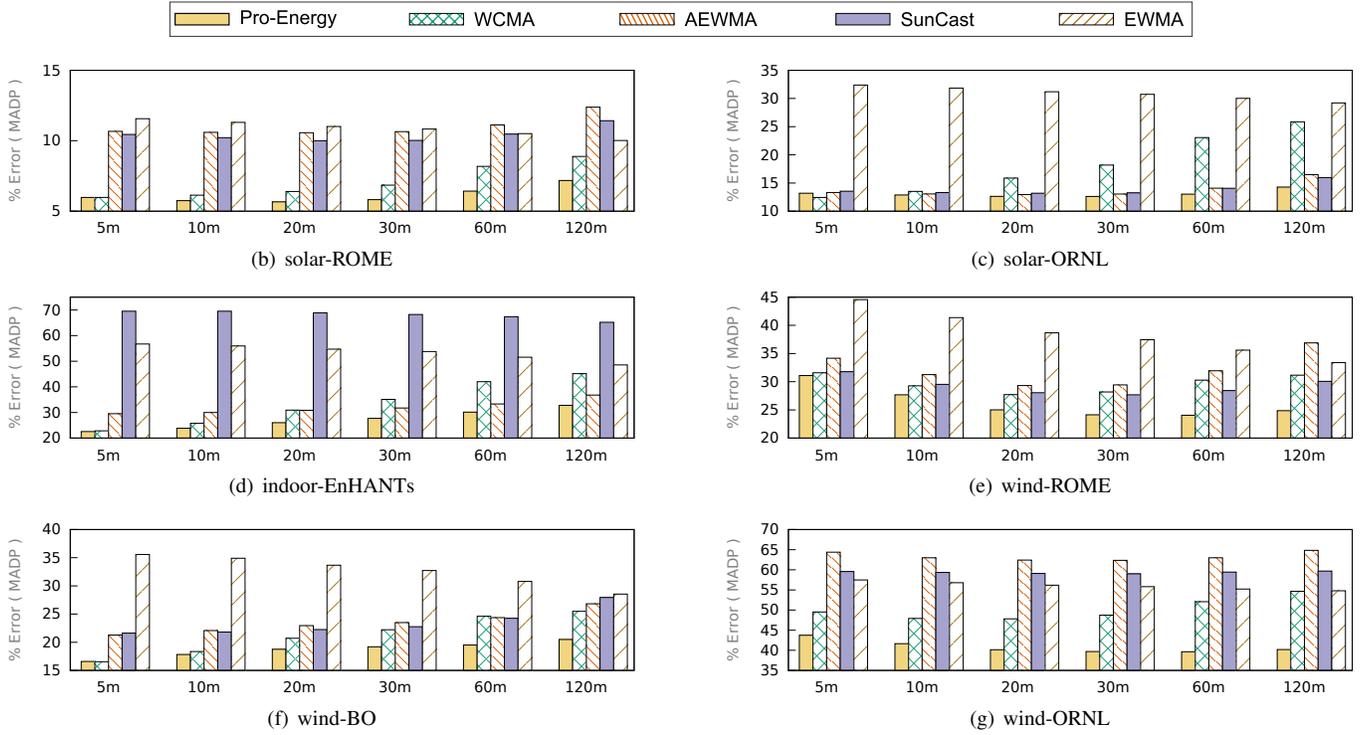


Fig. 6. Performance comparison of Pro-Energy against WCMA, AEWMA, SunCast and EWMA over different datasets.

C. Pro-Energy-VLT: Evaluation results

In this section we evaluate the impact of using variable-length timeslots on energy prediction accuracy by comparing the performance of Pro-Energy-VLT with that of Pro-Energy with fixed equal-length timeslots. In addition, we assess the performance of the online time-slotting technique proposed in Section IV by performing a comparative performance evaluation of Pro-Energy-VLT using two different algorithms for timeslots resizing and updating. These variants of Pro-Energy-VLT, which we termed Pro-Energy-VLT-adapt and Pro-Energy-VLT-PIP, are described in the following. In Pro-Energy-VLT-adapt, the distribution of the timeslots is computed according to the method proposed by Renner and Turau in [22]. To apply their algorithm, a fine-grained sampling of the harvesting process must be performed. As the minimum timeslot size we use in Pro-Energy-VLT is of 5 minutes, a fine-grained representation of the current day is kept by storing the amount of energy harvested over timeslots of 5 minutes. In Pro-Energy-VLT-PIP, the size of the N timeslots is determined by running an iterative algorithm based on the so-called Perceptually Important Point (PIP) method [35], an heuristic that identifies the $N + 1$ points with the greatest impact on the shape of the daily harvesting profile. As Pro-Energy-VLT-PIP requires a fine-grained representation of the energy harvesting profiles, in evaluating its performance the power source is sampled every 30 seconds and the results is stored in memory. Results are shown in Figure 7. The effectiveness of our proposed online time-slotting technique is confirmed by simulation results, which show how using Pro-Energy-VLT-adapt and Pro-Energy-VLT-PIP results in a MADP error that is, on average, 13.67% and 9.34% higher

than that of Pro-Energy-VLT, respectively. Results also confirm the benefit of using variable-length timeslots for both solar and wind energy predictions: Using Pro-Energy-VLT allows to significantly reduce the prediction error with respect to Pro-Energy. In case of solar energy predictions, combining Pro-Energy with variable-length timeslots further reduces the average MADP error of 10%-40% (Fig. 7(b)), 3%-22% (Fig. 7(c)) and 1%-25% (Fig. 7(d)), depending on the considered dataset. The average performance improvement is lower in case of wind energy predictions due to the higher variability of wind, which makes it more difficult to determine a timeslots distribution that well adapts to weather conditions over different days. Nevertheless, Pro-Energy-VLT achieves a MADP error of up to 2.59%, 6.54% and 1.21% lower than that Pro-Energy for wind-ROME, wind-BO and wind-ORNL datasets, respectively. As for energy consumption and memory overhead, Figure 8 shows the comparison between the error obtained by Pro-Energy-VLT for increasing values of N and the minimum prediction error of Pro-Energy, WCMA, AEWMA and SunCast, achieved when $N = 48$. The prediction error of EWMA is not reported, as it is always greater than 29%. Pro-Energy-VLT outperforms other models in terms of prediction error, using a number of timeslots that is between two fifth and two third of that used by Pro-Energy. In particular, for a prediction horizon of 10 minutes, Pro-Energy-VLT obtains the same accuracy of Pro-Energy using just 28 timeslots, and the same accuracy of WCMA, AEWMA and SunCast with a even smaller number of timeslots. In case of predictions at 60 minutes, Pro-Energy-VLT outperforms Pro-Energy using just 32 timeslots (instead of 48). This results in significant energy saving, as the overhead of an energy prediction model is directly proportional to N [34]. The

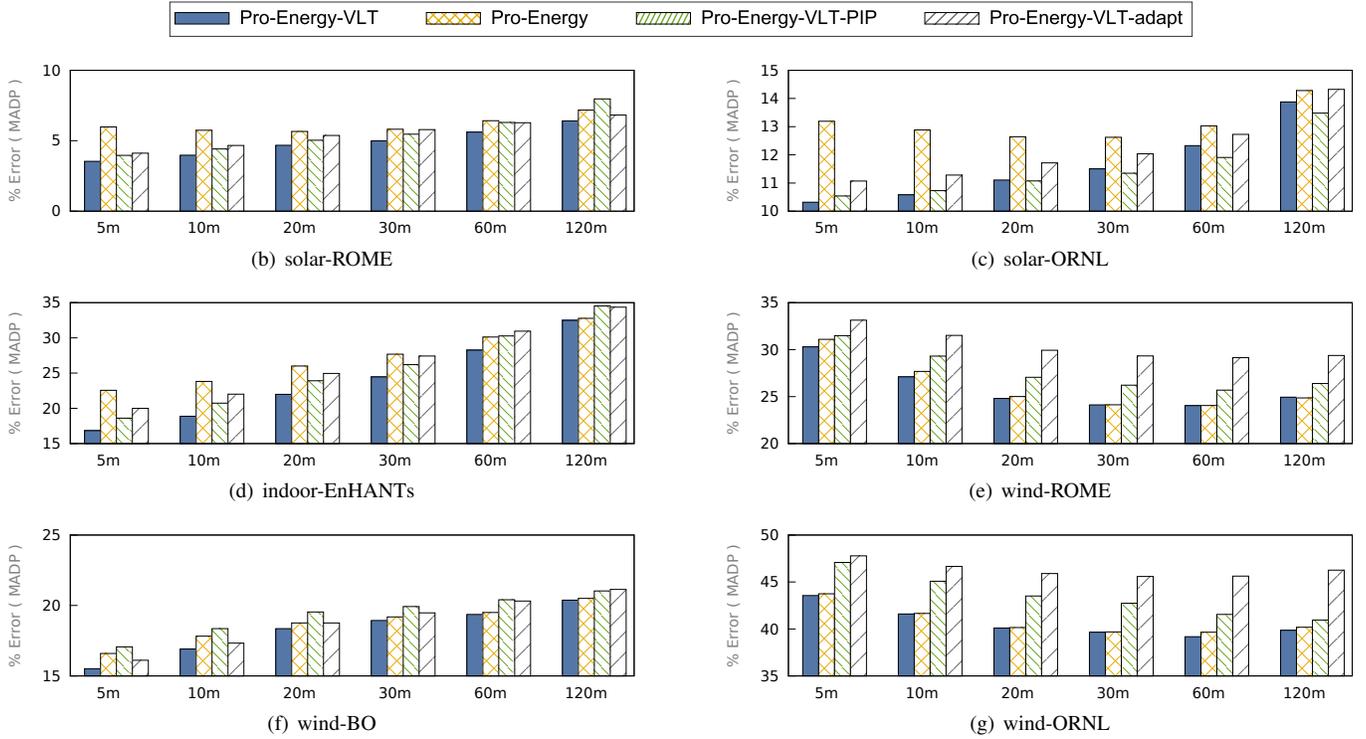


Fig. 7. Performance comparison of Pro-Energy-VLT against Pro-Energy, Pro-Energy-VLT-adapt and Pro-Energy-VLT-PIP over different datasets.

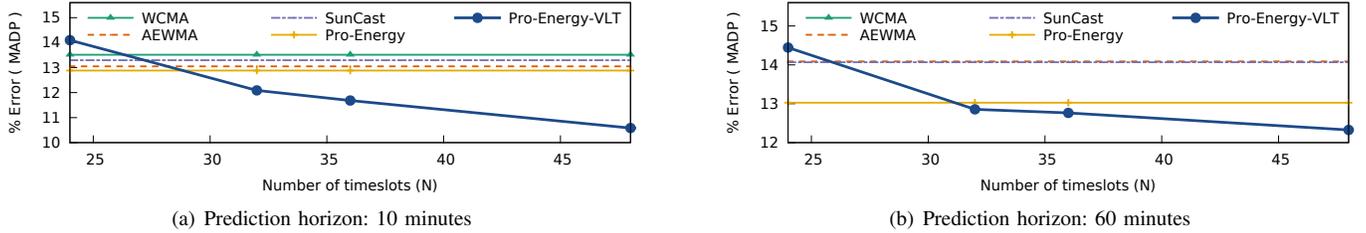


Fig. 8. Performance comparison of Pro-Energy-VLT with increasing values of N against competing schemes with $N=48$, for prediction horizons of a) 10 minutes and b) 60 minutes. The prediction error of WCMA is not reported in b) because it is greater than 23%.

details of the performance of each energy predictor in this scenario, in terms of average MADP error, average number of multiplications per day and memory overhead in bytes, is reported in Table II. The number of timeslots also affects the memory overhead of a predictor. For example, assuming harvesting samples are stored as 16 bit values, $N = 48$ and $D = 20$, WCMA and Pro-Energy require approximately 2 KB of RAM to store the matrix of the D previous days needed for energy predictions, which represents 20-50% of the total memory available on Telos B and Mica2/MicaZ motes, respectively. Using just 28 timeslots, Pro-Energy-VLT obtains a prediction error lower of that of WCMA and Pro-Energy, reducing the memory footprint of more than 40%.

D. Case study

To evaluate the performance improvement achieved by using energy prediction in a typical WSN scenario, we present a case study for task allocation supported by energy prediction. Finally, we assess the practical feasibility of our proposed approach through implementation on solar-powered motes.

TABLE II
PERFORMANCE OF ENERGY PREDICTORS: ERROR, AVERAGE NUMBER OF MULTIPLICATIONS PER DAY, MEMORY OVERHEAD IN BYTES

Predictor	% error	Multiplications	Memory
Pro-Energy-VLT (N=32)	12.67	793	1344
Pro-Energy (N=48)	12.88	1083	2592
WCMA (N=48)	13.51	335	1732
AEWMA (N=48)	13.05	3694	96
EWMA (N=48)	31.85	94	96

For the case study, we consider an application scenario where solar-powered nodes are deployed for periodically monitoring of environmental parameters. Nodes are also periodically requested to perform some energy-costly tasks that must be run to completion. We verify the performance of a *predictive* strategy that makes use of our energy prediction model to decide whether to accept or to refuse such incoming tasks. According to this strategy, a task is accepted only if its energy requirements can be met by using both the energy currently available to the node and the energy the node expects to harvest in the near future. In particular, whenever a task arrives, the node checks whether the sum of the energy

currently stored in its supercapacitor and of the energy it expects to harvest within the duration of the task is enough to run the task to completion. If this is the case, the node accepts the task, otherwise it rejects it to avoid wasting energy in running a task that would probably not be completed. We consider two other strategies for benchmarking: a *conservative* one, according to which a task is accepted only if the amount of energy currently available is enough to meet its energy requirements (thus not considering energy prediction) and an energy-unaware *greedy* strategy that always accepts tasks.

We perform a comparative performance evaluation of the three task allocation strategies using GreenCastalia. We use the default settings of GreenCastalia for channel and radio models. We accurately model energy consumption and harvesting by setting models parameters based on experimental measurements of prototype solar-powered nodes and by using real-life traces of solar availability. In particular, we consider Telos B motes powered by the harvesting system described in Section V. The energy consumption and harvesting parameters include energy efficiency of the harvesting board and experimental parameters extraction to model supercapacitors' leakage (as detailed in [33]), as well as power consumption of the Telos B in different states. We evaluate the performance of the three strategies in terms of number of tasks successfully completed, i.e., level of service the nodes provide, and percentage of time nodes are able to stay alive. In particular, a task is successfully completed if a node accepts to execute the task and it does not run out of energy before the task is completed. The percentage of time nodes are able to stay alive, which we term availability, is the ratio of the number of seconds for which the supercapacitor voltage is greater than the minimum operating voltage of the node to the total number of seconds in the experiment.

Figure 9 shows a performance comparison of the three strategies. Different scenarios of energy availability are simulated by varying the size of the solar cell powering the nodes, and thus the amount of energy they harvest (Eq. 11). As expected, the predictive strategies offers the best tradeoff between achieved level of service (number of tasks successfully completed) and availability (in terms of percentage of time nodes in the network are alive). The conservative strategy achieves availability very close to 100%, but at the price of refusing many tasks that cannot be completed by solely using the energy available in the node supercapacitors. The predictive strategy, instead, accepts and successfully executes such tasks when environmental energy is available, thanks to the fact that the future energy intake is estimated using the prediction model. On average, the predictive strategy successfully executes 50% more tasks than the conservative one. The greedy strategy executes almost 10% more tasks than the predictive one, but its availability is low, especially for smaller solar cell sizes. In particular, if the size of the solar cell is set to 8 cm² or lower, the average dead time of the nodes is greater than 40%, which hinders the general functionality of the network (e.g., periodic sensing). Using the predictive strategy reduces the dead time of the nodes by 35% in case of large solar cells (i.e., solar cell size of 32 cm²) and up to a factor 23x in case a smaller solar cell is used.

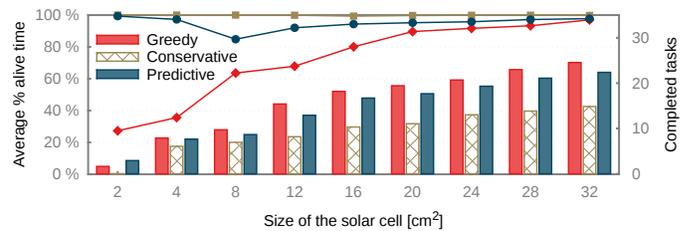


Fig. 9. Completed tasks and percentage of alive time for different allocation strategies and solar cell sizes.

Finally, to quantify the energy overhead introduced by energy prediction, we implemented the three task allocation strategies, together with Pro-Energy and Pro-Energy-VLT, in TinyOS [29]. Based on experimental measurements of power consumption, the additional energy overhead experienced by nodes actually running our proposed forecast algorithm is around 3 J *per day* when $N = 48$ with respect to the case in which no energy prediction algorithm is used.

VI. CONCLUSIONS

In this paper, we have presented Pro-Energy, a novel energy prediction model for multi-source energy harvesting WSNs. Pro-Energy stores and maintains a pool of *typical harvested profiles* observed in the past, which are combined with current observations to compute energy prediction at future timeslots for both short and medium prediction horizons. An extension of Pro-Energy, called Pro-Energy-VLT, is proposed to combine our energy prediction model with timeslots of variable lengths. Pro-Energy-VLT adapts the granularity of the prediction timeslots to the dynamics of the energy source, further improving prediction accuracy and reducing the memory footprint and the energy overhead of energy prediction. We have performed extensive validation of Pro-Energy and Pro-Energy-VLT using real-life traces of the harvested energy we have obtained by interfacing Telos B nodes with photovoltaic panels and wind micro-turbines, as well as publicly-available traces of solar and wind availability. Results show that Pro-Energy outperforms state-of-the-art energy predictors, providing improvements in prediction accuracy as high as 67%. With respect to competing energy prediction models, which achieves good prediction accuracy at the cost of high memory footprint, Pro-Energy-VLT obtains better prediction performance, while reducing the memory overhead of up to 40%.

REFERENCES

- [1] S. Basagni, M. Y. Naderi, C. Petrioli, and D. Spenza, "Wireless Sensor Networks with Energy Harvesting," in *Mobile Ad Hoc Networking: The Cutting Edge Directions*. Hoboken, NJ: John Wiley and Sons, Inc., Mar 2013, ch. 20, pp. 701–736.
- [2] D. Mishra, S. De, S. Jana, S. Basagni, K. Chowdhury, and W. Heinzelman, "Smart RF energy harvesting communications: challenges and opportunities," *IEEE Communications Magazine*, vol. 53, no. 4, pp. 70–78, April 2015.
- [3] T. Zhu, Z. Zhong, Y. Gu, T. He, and Z. L. Zhang, "Leakage-aware energy synchronization for wireless sensor networks," in *Proceedings of ACM MobiSys 2009*, Krakow, Poland, Jun 2009, pp. 319–332.
- [4] J. Marañević, C. Stein, and G. Zussman, "Max-min fair rate allocation and routing in energy harvesting networks: Algorithmic analysis," in *Proceedings of ACM MobiHoc 2014*, Philadelphia, PA, USA, Aug 2014, pp. 367–376.

- [5] N. Tekbiyik, T. Girici, E. Uysal-Biyikoglu, and K. Leblebicioglu, "Proportional fair resource allocation on an energy harvesting downlink," *IEEE Transactions on Wireless Communications*, vol. 12, no. 4, pp. 1699–1711, Apr 2013.
- [6] C. Moser, J.-J. Chen, and L. Thiele, "Power management in energy harvesting embedded systems with discrete service levels," in *Proceedings of ACM/IEEE ISLPED 2009*, San Francisco, CA, USA, Aug 2009, pp. 413–418.
- [7] X. Ren, W. Liang, and W. Xu, "Quality-aware target coverage in energy harvesting sensor networks," *IEEE Transactions on Emerging Topics in Computing*, vol. 3, no. 1, pp. 8–21, Mar 2015.
- [8] T. N. Le, M. Magno, A. Pegatoquet, O. Berder, O. Sentieys, and E. Popovici, "Ultra Low Power Asynchronous MAC Protocol Using Wake-up Radio for Energy Neutral WSN," in *Proceedings of ACM ENSSys 2013*, Rome, Italy, Nov 2013, pp. 10:1–10:6.
- [9] C. Renner, S. Unterschütz, V. Turau, and K. Römer, "Perpetual data collection with energy-harvesting sensor networks," *ACM Transactions on Sensor Networks*, vol. 11, no. 1, pp. 12:1–12:45, Sep. 2014.
- [10] T. N. Le, A. Pegatoquet, O. Sentieys, O. Berder, and C. Belleudy, "Duty-cycle power manager for thermal-powered Wireless Sensor Networks," in *Proceedings of IEEE PIMRC 2013*, London, UK, Sept 2013, pp. 1645–1649.
- [11] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava, "Power management in energy harvesting sensor networks," *ACM Transactions on Embedded Computing Systems*, vol. 6, no. 4, Sep 2007.
- [12] C. Bergonzini, D. Brunelli, and L. Benini, "Comparison of energy intake prediction algorithms for systems powered by photovoltaic harvesters," *Elsevier Microelectronics Journal*, vol. 41, no. 11, pp. 766–777, Nov. 2010.
- [13] J. Piorno, C. Bergonzini, D. Atienza, and T. Rosing, "Prediction and management in energy harvested wireless sensor nodes," in *Proceedings of Wireless VITAE 2009*, Aalborg, Denmark, May 2009, pp. 6–10.
- [14] D. K. Noh and K. Kang, "Balanced energy allocation scheme for a solar-powered sensor system and its effects on network-wide performance," *Elsevier Journal of Computer and System Sciences*, vol. 77, no. 5, pp. 917–932, Sep 2011.
- [15] J. Lu, S. Liu, Q. Wu, and Q. Qiu, "Accurate modeling and prediction of energy availability in energy harvesting real-time embedded systems," in *Proceedings of IEEE IGCC 2010*, Chicago, IL, USA, Aug 2010, pp. 469–476.
- [16] C. Renner, "Solar harvest prediction supported by cloud cover forecasts," in *Proceedings of ACM ENSSys 2013*, Rome, Italy, Nov 2013, pp. 1:1–1:6.
- [17] C. Renner and P. Nguyen, "A system for efficient dissemination of weather forecasts for sustainable solar-powered sensors," in *Proceedings of IFIP SustainIT 2015*, Madrid, Spain, Apr 2015, pp. 1–8.
- [18] C. Renner and P. A. T. Nguyen, "Lossless Compression of Cloud-cover Forecasts for Low-overhead Distribution in Solar-harvesting Sensor Networks," in *Proceedings of ACM ENSSys 2014*, Memphis, TN, USA, Nov 2014, pp. 43–48.
- [19] N. Sharma, J. Gummeson, D. Irwin, T. Zhu, and P. Shenoy, "Leveraging weather forecasts in renewable energy systems," *Sustainable Computing: Informatics and Systems*, vol. 4, no. 3, pp. 160–171, 2014.
- [20] J. Lu and K. Whitehouse, "SunCast: Fine-grained Prediction of Natural Sunlight Levels for Improved Daylight Harvesting," in *Proceedings of ACM IPSN 2012*, Beijing, China, Apr 2012, pp. 245–256.
- [21] N. Sharma, P. Sharma, D. Irwin, and P. Shenoy, "Predicting solar generation from weather forecasts using machine learning," in *Proceedings of SmartGridComm 2011*, Brussels, Belgium, October 2011, pp. 528–533.
- [22] C. Renner and V. Turau, "Adaptive energy-harvest profiling to enhance depletion-safe operation and efficient task scheduling," *Sustainable Computing: Informatics and Systems*, vol. 2, no. 1, pp. 43–56, 2012.
- [23] A. Cammarano, C. Petrioli, and D. Spenza, "Pro-Energy: a novel energy prediction model for solar and wind energy harvesting WSNs," in *Proceedings of IEEE MASS 2012*, Las Vegas, NV, USA, Oct 2012, pp. 75–83.
- [24] A. Cammarano, C. Petrioli, and D. Spenza, "Improving Energy Predictions in EH-WSNs with Pro-Energy-VLT," in *Proceedings of ACM SenSys 2013, Poster Session*, Rome, Italy, Nov 2013, pp. 41:1–41:2.
- [25] IXYS Corporation, "XOB17 IXOLAR High Efficiency Solar Bits technical information," Jun 2009.
- [26] M. Gorlatova, A. Wallwater, and G. Zussman, "Networking low-power energy harvesting devices: Measurements and algorithms," in *Proceedings of IEEE INFOCOM 2011*, Shanghai, China, Apr 2011, pp. 1602–1610.
- [27] "NREL: Measurement and Instrumentation Data Center," 2011, <http://www.nrel.gov/mid/>.
- [28] D. Carli, D. Brunelli, L. Benini, and M. Ruggeri, "An effective multi-source energy harvester for low power applications," in *Proceedings of IEEE DATE 2011*, March 2011, pp. 1–6.
- [29] P. Levis, S. Madden, J. Polastre, R. Szewczyk, K. Whitehouse, A. Woo, D. Gay, J. Hill, M. Welsh, E. Brewer, and D. Culler, "TinyOS: An Operating System for Sensor Networks," in *Ambient Intelligence*. Springer Berlin Heidelberg, 2005, pp. 115–148.
- [30] J. Azevedo and F. Santos, "Energy harvesting from wind and water for autonomous wireless sensor nodes," *IET Circuits, Devices & Systems*, vol. 6, no. 6, pp. 413–420, Nov 2012.
- [31] D. Benedetti, C. Petrioli, and D. Spenza, "GreenCastalia: An energy-harvesting-enabled framework for the Castalia simulator," in *Proceedings of ACM ENSSys 2013*, Rome, Italy, November 14 2013, pp. 7:1–7:6.
- [32] A. Boulis, "Castalia: Revealing Pitfalls in Designing Distributed Algorithms in WSN," in *Proceedings of ACM SenSys 2007*, 2007, pp. 407–408.
- [33] T. L. Porta, C. Petrioli, C. Phillips, and D. Spenza, "Sensor Mission Assignment in Rechargeable Wireless Sensor Networks," *ACM Transactions on Sensor Networks*, vol. 10, no. 4, pp. 60:1–60:39, Jun. 2014.
- [34] M. Ali, B. Al-Hashimi, J. Recas, and D. Atienza, "Evaluation and design exploration of solar harvested-energy prediction algorithm," in *Proceedings of IEEE DATE 2010*, Dresden, Germany, Mar 2010, pp. 142–147.
- [35] F. Chung, T. Fu, R. Luk, and V. Ng, "Flexible time series pattern matching based on perceptually important points," in *Proceedings of IJCAI 2001, Workshop on Learning from Temporal and Spatial Data*, Seattle, Washington, Aug 2001, pp. 1–7.



Alessandro Cammarano received the Ph.D. in Computer Science from University of Rome "La Sapienza" in 2014. His research interests focus on wireless sensor networks design, cognitive radio networks, energy harvesting wireless sensor networks and energy prediction models. Dr. Cammarano has been a member of La Sapienza team within the FP7 project SENDORA, in which he has designed, evaluated and simulated a complete solution for wireless sensor networks for spectrum monitoring.



Chiara Petrioli received the Ph.D. in Computer Engineering from University of Rome "La Sapienza" in 1998. She is currently a full professor at University of Rome "La Sapienza", Computer Science Department, where she leads the Sensor Networks and Embedded Systems laboratory (SENSES lab). She also leads the Cyber Physical System lab of "La Sapienza" center for Cyber Intelligence and Information Security, and is a founding partner of "La Sapienza" spinoff WSENSE S.r.l. Her research interests focus on the design and optimization of wireless, embedded and cyber physical systems. Prof. Petrioli is member of the steering committee of ACM SenSys, IEEE SECON and Mobiquitous, program co-chair of IEEE INFOCOM 2016, workshop co-chair of ACM MobiCom 2014, general chair of ACM SenSys 2013. She has been member of the steering committee and associate editor of IEEE Transactions on Mobile Computing, associate editor of IEEE Transactions on Vehicular Technology, member of the executive committee of ACM SIGMOBILE, and has been program co-chair of leading conferences in the field, such as ACM MobiCom and IEEE SECON. Prof. Petrioli has published over a hundred papers in prominent international journals and conferences. She has also been PI of over twenty national and international research projects, serving as coordinator of two EC projects (FP7 projects GENESI and SUNRISE). Prof. Petrioli was a Fulbright scholar.



Dora Spenza is a Postdoc researcher at the Department of Computer Science of University of Rome "La Sapienza". She received the Ph.D. in Computer Science from the same university in 2013. In 2011, she was a visiting scholar at the Computer Science and Engineering Department at Penn State University. Her research interests concern the design and optimization of wireless embedded sensing systems, with an emphasis on ultra-low-power and energy-harvesting techniques for very long-lasting wireless sensor networks. She is the recipient of the 2010

Google European Doctoral Fellowship in Wireless Networking.