Sensor-mission Assignment in Wireless Sensor Networks with Energy Harvesting

Tom La Porta*, Chiara Petrioli† and Dora Spenza†,
*Department of Computer Science and Engineering, Pennsylvania State University, USA
†Department of Computer Science, Sapienza University of Rome, Italy
E-mail: tlp@cse.psu.edu, {petrioli,spenza}@di.uniroma1.it

***Abstract*—Sensor mission assignment concerns matching the sensing resources of a wireless sensor network (WSN) to appropriate tasks (missions), which may come to the network dynamically. Although solutions for WSNs with battery-operated nodes have been proposed for this problem, no attention has been given to networks whose nodes have energy harvesting capabilities, which impose quite a different energy model. In this paper we address this problem by providing both an analytical model and a distributed heuristic, called EN-MASSE, for energy harvesting WSNs. The objective of both model and EN-MASSE is to maximize the profit of the network, fully exploiting the harvesting technologies, while ensuring the execution of the most critical missions within a given target WSN lifetime. The performance of EN-MASSE is evaluated by simulations based on real solar energy traces. Our experiments show that EN-MASSE behaves very closely to the optimum provided by our model and significantly outperforms previously proposed solutions.**

I. Introduction

Field and environmental monitoring are very common applications of wireless sensor networks (WSNs). In a typical scenario, the nodes of a WSN are placed in a region for collecting accurate measurements about some phenomenon. The number of active nodes in the network (i.e., those currently performing their sensing and communication tasks) determines the accuracy of the collected data and the number of tasks that can be complete successfully. There exists a tradeoff between the accuracy and number of tasks that may be performed, and the longevity of the network. Since WSNs are usually intended to last for long periods of time, the number of active nodes should be carefully managed.

Recent works [1] have addressed selective activation of nodes, i.e., sensing tasks are allocated only to a group of sensor nodes over time, allowing the other nodes in the network to go to sleep and save energy. In many applications, however, the application may require the network to achieve multiple, simultaneous missions, potentially competing for the sensing resources of the nodes. By *mission*, we refer to a sensing task whose primary goal is the collection of information, to which one or multiple sensors may contribute. For example, the network might be required to perform multiple, concurrent localization and intrusion detection tasks using directional acoustic sensors that can be assigned to only one mission at time [2]. In this case, a sensor selection scheme is no longer sufficient, because it is necessary to decide not only which nodes in the network must be active, but also which active sensors should be assigned to which mission. This is a non trivial task, because a given node may offer support to different missions with different levels of accuracy and fit. Missions, on the other hand, may vary in importance (*profit*) and amount of resources they require (*demand*). They may also appear in the network at any time and may have different durations. A mission is executed if enough sensors are assigned to it, based on its demand. In this scenario, the goal of a sensor-missions assignment algorithm is thus to assign available nodes to appropriate missions, maximizing the profit received by the network for mission execution.

Different constraints in assignment decision lead to different versions of this problem, depending on whether missions can be assigned to sensors dynamically [2]–[4] or they are known a priori [3]–[5] and on whether the network has a predefined target operational lifetime [4] or must last as long as possible [3]. Some prior solutions are energy aware, in the sense that they take into account nodal residual energy to decide mission assignments. In doing so, however, they make the specific assumption that energy is monotonically decreasing, as is typical with a battery, and therefore residual energy is the only criterion for assigning missions.

More recently, sensor nodes are being developed that, along traditional batteries, mount devices such as *energy harvesters* and *supercapacitors*, which draw new energy from the environment surrounding a node and store it, respectively. For these nodes, new paradigms for mission assignments are needed, which take into account that nodes currently having little or no energy left might have enough in the future to carry out new missions.

In this paper we are concerned with the practical case of WSNs that run applications requiring the network to be operational for a given amount of time (*dynamic assignment with a time horizon*, as in [4]). We study this case with sensor nodes that have energy harvesting capabilities. We refer to the given amount of time the network is expecting to be operational, i.e., accepting and completing missions, as the *target lifetime* of the network.

We aim at making the following contributions:

1) We model the problem of optimum mission assignment in energy harvesting WSNs. Our formulation captures the details of the behavior of a typical energy harvesting subsystem, selects which missions to execute and the nodes that should execute them so that the profit throughout the network operational lifetime is maximized. At the same time, the schedule

978-1-4577-0093-4/11/$26.00

provided by our formulation ensures that the required energy is always available, independently of its source (battery or harvesting).

2) We provide the first ENergy harvesting-aware sensor-Mission ASSignmEnt distributed algorithm (denoted EN-MASSE). In order to decide whether to bid for a given mission or not, an EN-MASSE node exploits not only the nodal residual energy, but also takes into account the energy that is expected to be harvested in the future (based on a harvester-based energy prediction model) and the expected profit and demand of missions to come. Since our approach is general, EN-MASSE may be combined with any energy prediction model.

3) We provide a simulation based performance evaluation framework for energy harvesting WSNs. In our experiments, we use traces of the availability of solar energy that we obtained by interfacing TelosB nodes with solar cells, collecting data for 100 days.

4) We perform a comparative performance evaluation of EN-MASSE, the schemes proposed by Johnson et al. in [4], as well as the optimum mission allocation computed by our analytical model. We evaluate the impact of critical parameters, such as the target lifetime, the types of sensors embedded in the nodes (and their energy consumption), the supercapacitor size and the mission arrival rate, on the performance of the different schemes in twenty distinct scenarios. Our results show that the profit earned by EN-MASSE is close to the optimum. EN-MASSE is also able to significantly outperform previous schemes in terms of profit, of efficient usage of the harvested energy and of the capability to support critical missions over time.

The remainder of this paper is organized as follows. Related work is presented in Section II. In Section III we formulate the mission assignment problem for energy harvesting WSNs and provide an analytical model that maximizes network profit. EN-MASSE is then presented in Section IV and compared to previous solutions and the optimum derived by the model in Section V. Finally, we present our conclusions in Section VI.

II. Related work

A. Sensor-mission assignment problem

Mission assignment in WSNs has received considerable attention. Bar-Noy et al. introduced the Semi-Matching with Demands (SMD) in [5]. This approach is based on different priorities and demands of each mission and on additive utility values for each sensor-mission pair. In the original SMD profits are awarded only if a certain utility threshold is met and the problem is defined only for a set of missions known a priori. SMD was extended in [3], incorporating both a profit threshold in case of partial mission satisfaction and mission dynamics. The authors propose centralized and distributed approaches for maximizing the network profit by satisfying all missions available at a given time. They also provide an energy-aware assignment scheme for prolonging the network lifetime. However, the case in which the network operates for a finite target lifetime is not considered.

A variation of the sensor-mission assignment problem, motivated by frugality and conservation of resources, was addressed by Johnson et al. in [4]. The authors show that finding an optimal solution to the dynamic sensor-mission assignment problem is NP-hard. They then propose a heuristic where the assignment decisions depend on the sensors energy, exposing a trade-off between network lifetime and achievable profit. Rather than satisfying all missions available at a given time, this scheme allows the nodes to autonomously decide the missions in which they will participate based on their residual energy.

Other related problems with different assignment constraints have been studied. For example, in [6], Le et al. consider missions than can be decomposed into a set of specific tasks and solve the sub-problem of sensor-task assignment, allowing sensors to be shared and reassigned between mission subtasks.

B. Power management in energy harvesting networks

Despite significant research effort, energy continues to remain a severe bottleneck for applications where battery-powered systems are expected to operate for long periods of time. For this reason there has been a growing interest in the design of systems that are able to draw energy from the environment, with the main goal of supplementing or even replacing batteries (*energy harvesters*) [7]–[10]. Due to the great variability of environmental energy sources and to their unpredictable nature, harvesting-aware power management policies are required for network performance and lifetime enhancement. Many works have addressed the possibility of achieving near perpetual network lifetime by operating in *energy-neutral* mode, i.e., by consuming only as much energy as harvested [11], [12].

Other works have tackled the problem of task scheduling for energy harvesting WSNs [13]–[15]. For example, in [15], Steck and Rosing presented two algorithms to balance the tradeoff between task utility and energy constraints in these networks, guaranteeing energy neutrality.

Our approach differs from these works in two ways. First, our primary goal is to maximize the network profit within a given time horizon, rather than enabling the network to operate perennially. We thus do not require energy neutrality. Second, we focus on the problem of assigning nodes in the network to competitive missions, instead of scheduling local tasks of nodes.

III. Energy Harvesting Architecture and Analytical Model

Each node in the network is equipped with an energy harvesting subsystem, which includes one or more photovoltaic panels[1] and a supercapacitor that acts as energy buffer for the node . Nodes are also equipped with a non-rechargeable primary battery of capacity PB, whose purpose is to guarantee a minimum lifetime in case no or little energy can be drawn from the environment.

[1]Although we focus here on solar-based harvesters, our approach is general and can be applied to other energy sources as well.

An actual supercapacitor deviates from an ideal energy buffer in many ways. First of all, it has a finite size B^{Max} and therefore can hold a finite amount of energy. Second, it suffers from leakage and self-discharge, through which energy is lost even if the supercapacitor is not in use. Finally, it has a charging efficiency $\eta_c < 1$ and a discharging efficiency $\eta_d < 1$, i.e., some energy is lost while charging and discharging the supercapacitor.

The leakage experienced by a charged supercapacitor is a complex function, that increases when the energy stored is higher [7], [16]. To capture this behavior, we model the leakage, $leak_i(t)$, experienced by the energy buffer, B_i, of the sensor node N_i at time t by using a piecewise linear approximation of the empirical leakage pattern, as in [16]:

$$leak_i(t) = \begin{cases} a_1 \cdot B_i(t) + b_1, & B_{R_1} \leq B_i(t) < B_{R_2} \\ \vdots & \vdots \\ a_n \cdot B_i(t) + b_n, & B_{R_n} \leq B_i(t) < B_{R_{n+1}} \end{cases}$$

where $B_{R_1}, \ldots, B_{R_{n+1}}$ are the residual energy values in which the slope of the leakage curve change significantly and $a_1, \ldots, a_n, b_1, \ldots, b_n$ are constants depending on the supercapacitor used, that represent the coefficients of the line segments used for the approximation.

The mission assignment problem for wireless sensor networks equipped with such an energy harvesting subsystem is then mathematically formulated as follows. The network consists of a set of energy-harvesting endowed sensor nodes $N_1, \ldots, N_n$, pre-deployed in a field. At any time, a mission may appear in the network at a specific geographic location. Let $M_1, \ldots, M_m$ be the set of missions that the network is asked to perform.

A mission M_j is a tuple $(p_j, d_j, l_i, ts_i, te_i)$ where:

- p_j is the profit of the mission, indicating both its importance and the reward achieved by the network for its execution;
- d_j is the mission demand, indicating the amount of sensing resources it needs;
- l_i is the geographic location of the mission in the field;
- ts_i is the time the mission arrives in the network;
- te_i is the time the mission terminates.

Missions usually last for multiple units of time, i.e., $te_j > ts_j$; p_j, d_j, l_j are constant over time during mission execution.

In order for a mission to be executed, one or more sensor nodes must be assigned to it. We define e_{ij} as the utility received by mission M_j if node N_i is assigned to it. This utility value is zero if the node cannot contribute to a mission, which happens, for instance, if the mission location and the node position are not close enough. Different values of e_{ij} can be used to indicate the "quality of contribution" that a sensor can provide to a particular mission.

While missions can be performed by multiple sensors, we assume that a node can be assigned to at most one mission at a time ($x_{ij} = 1$). We also assume that the total utility received by a mission u_j is equal to the sum of the utilities provided by the sensors assigned to it, i.e., $u_j = \sum x_{ij} e_{ij}$.

The total utility that a mission M_j requires is expressed by its demand d_j. If the total utility u_j received by a mission is lower than its requested demand d_j, we say that the mission M_j is partially satisfied and we indicate its satisfaction level with $y_j = u_j/d_j$ (in the range $[0, 1]$). Profits are received by the network for mission execution, based on the satisfaction level of the mission. In this formulation of the mission assignment problem, profits can be awarded fractionally, but only if a minimum satisfaction threshold T_{sat} is met.

More formally, the profit associated with satisfying mission M_j at time t depends on the satisfaction level y_{jt} of the mission at time t, as follows:

$$p_{jt}(y_{jt}) = \begin{cases} p_j, & \text{if } y_{jt} \geq 1 \\ p_j \cdot y_{jt}, & \text{if } T_{sat} \leq y_{jt} < 1 \\ 0 & \text{if } y_{jt} < T_{sat} \end{cases}$$

Since missions usually last for multiple units of time, the total profit received for achieving mission M_j is the sum of the profits earned over the entire mission lifetime, i.e., $p_j = \sum_{t=ts_j}^{te_j} p_{jt}(y_{jt})$.

When a node is assigned to a mission, it spends a certain amount of energy per unit time, sc to accomplish the sensing activity required by the mission. When a sensor is not assigned to any mission, its power consumption is equal to the idle energy consumption, ic. Thus, the energy consumption $EC_i(t)$ of the sensor node N_i at time t is:

$$EC_i(t) = \begin{cases} 0, & \text{if the node is dead} \\ sc, & \text{if } \exists M_j \in M \text{ s.t. } x_{ijt} = 1 \\ ic, & \text{if } \forall M_j \in M, x_{ijt} = 0 \end{cases}$$

where we consider a node *dead* at time t, and thus set its energy consumption to zero, if it has not enough energy to stay in idle mode.

The amount of energy a node obtains from the harvesting subsystem is modeled taking into account features of realistic energy buffers. We consider two different cases, based on the amount of energy harvested at a certain time t, the current energy consumption and the energy level of the supercapacitor: 1) If the current energy consumption is greater than (or equal to) the energy currently harvested, then the node can directly use the harvested energy to (partially) fulfill its power requirements. This is the most efficient way of using the environmental energy, because there is no energy loss due to buffer inefficiency and leakage. 2) Otherwise, some energy is directly used to sustain the node's operation, while excess energy is stored in the supercapacitor for later use.

Because both the charging and discharging efficiency of the buffer are strictly less than one, only a fraction $\eta_c \eta_d$ of the excess energy is available after storing (and retrieving) it. Furthermore, because of the finite size of the buffer, some energy may be wasted if there is not enough space left in the supercapacitor to store it. More formally, the energy obtained from the harvesting subsystem, $EH_{ij}(t)$, and *directly* used by the sensor node N_i at time t is:

$$EH_i(t) = \begin{cases} ES_i(t), & \text{if } ES_i(t) \leq EC_i(t) \\ EC_i(t), & \text{if } ES_i(t) > EC_i(t) \end{cases}$$

where $ES_i(t)$ is the energy harvested through the solar panel by the node N_i at time t.

The energy provided or stored by the supercapacitor at time t is instead:

$$EB_i(t) = \begin{cases} \eta_d(EC_i(t) - ES_i(t)), & \text{if } ES_i(t) \leq EC_i(t) \\ \eta_c(ES_i(t) - EC_i(t)), & \text{if } ES_i(t) > EC_i(t) \end{cases}$$

In the first case, the supercapacitor provides enough energy to sustain the current consumption of the node and is discharged with a discharging efficiency equal to η_d. In the second case, excess energy generated by the harvesting subsystem is stored in the supercapacitor for later use with a charging efficiency of η_c.

The difference $\Delta B_i(t)$ in the supercapacitor energy between time t and $t+1$ can then be computed as:

$$\Delta B_i(t) = \begin{cases} EB_i(t), & \text{if } ES_i(t) > EC_i(t) \text{ and} \\ & EB_i(t) \leq B_i^{Max} - B_i(t) \\ B_i^{Max} - B_i(t), & \text{if } ES_i(t) > EC_i(t) \text{ and} \\ & EB_i(t) > B_i^{Max} - B_i(t) \\ -EB_i(t), & \text{if } ES_i(t) \leq EC_i(t) \text{ and} \\ & EB_i(t) \leq B_i(t) \\ -B_i(t), & \text{if } ES_i(t) \leq EC_i(t) \text{ and} \\ & EB_i(t) > B_i(t) \end{cases}$$

The energy stored in the supercapacitor at time $t+1$, $B_i(t+1)$ is finally given by: $B_i(t+1) = B_i(t) + \Delta B_i(t) - leak_i(t)$.

Sensor-mission assignment problem. Ideally, we seek an assignment of sensors to missions that satisfies each mission's demand. However, satisfying all missions may not be feasible, thus our goal is to maximize the total profit obtained by the network over a given target lifetime. The sensor-mission assignment problem in energy harvesting WSNs is modeled as follows:

$$\begin{aligned} &\max \textstyle\sum_t^{T_l} \sum_{j=1}^m p_j(y_{jt}) && && (1)\\ &\text{s.t.} \textstyle\sum_{i=1}^n x_{ijt} e_{ij} \geq d_j y_{jt}, && \forall M_j, t && (2)\\ &\textstyle\sum_{j=1}^m x_{ijt} \leq 1, && \forall N_i, t && (3)\\ &\textstyle\sum_{t' \leq t} EC_i(t') \leq PB_i + && && \\ &\quad \textstyle\sum_{t' \leq t} (EH_i(t') + B_i(t')) && \forall N_i, t && (4)\\ &x_{ijt} \in \{0, 1\}, && \forall x_{ijt} && (5)\\ &y_{jt} \in [0, 1], && \forall y_{jt} && (6)\\ &B_i(0) = 0, && \forall N_i && (7) \end{aligned}$$

where t varies between 0 and the network target lifetime T_l.

There are two sets of decision variables: y_{jt} indicating the satisfaction level of the mission M_j at time t and x_{ijt}, denoting if sensor N_i is assigned or not to mission M_j at time t. We seek an assignment of sensors to missions that maximizes the total profit obtained by the network over a given target lifetime T_l (1). The satisfaction level y_{jt} of the mission M_j at time t depends on the utility received by the mission with respect to its demand (2). A node may be assigned to at most one mission at a time (3). For each sensor node N_i and for each time instant t, we ask that the total energy consumed by N_i in the first t slots is always equal or less than the energy initially stored in its primary battery[2], plus the sum of the harvested energy directly used and the harvested energy obtained through the supercapacitor within the same interval of t slots. Using this condition, and having accurately modeled the harvesting subsystem behavior, we assure that the energy constraints are not violated at any instant of time. Finally, nodes can not be fractionally assigned to a mission (5), the satisfaction level of a mission is in the range $[0, 1]$ (6) and the supercapacitor is initially empty (7).

IV. EN-MASSE

Each mission appears in the network at a specific geographic location l_i. In EN-MASSE the sensor node closest to l_i is selected as the *mission leader* and coordinates the process of assigning nodes to the mission. Nodes in the neighborhood of the mission decide autonomously whether to bid for the execution of the incoming mission or to ignore it, based on a *proposal (or bidding) scheme*. They then communicate their possible availability to the mission leader, which greedily selects which of the available sensor nodes to assign to the mission, based on their offered contribution, until either the mission is fully satisfied or all nodes which have bid for it have been assigned to the mission.[3] We selected the communication protocol described in [3], [4] for exchanging information between the mission leader and the nearby nodes.

The EN-MASSE bidding scheme is specific for sensor-mission assignment in networks with energy harvesting capabilities. Each time a sensor becomes aware of a mission, it considers several factors in order to decide whether to propose for its execution or not, including: 1) the current energy level of the node battery and capacitor; 2) the energy cost of the mission; 3) the future energy availability, obtained through a solar energy prediction model; 4) the profit of the mission with respect to the maximum profit; 5) the utility offered by the node with respect to the mission demand, and 6) the target lifetime of the network.

The first three factors in the list, namely the current energy level of the battery and the supercapacitor, the energy cost of the mission and the expected future energy available, are used by the node to classify the incoming mission into one of the following four classifications: **Free missions** are those arriving when the node super-capacitor is full; their energy cost is expected to be fully sustained by the energy harvested during their duration. The information on the energy that will be harvested in the near future is provided by the energy prediction model. **Recoverable missions** are those whose energy cost can be sustained using the energy stored in the supercapacitor. Such energy cost can be recovered through harvesting in a small period of time, according to the prediction of future energy availability. **Capacitor-sustainable missions** occur if the total energy cost of the mission can be

[2] The size of the battery must be chosen as to guarantee a node being idle all the time to reach at least the target lifetime, even if the energy source is unavailable over the whole period.

[3] Sensor nodes are actually allocated to missions only if such missions can be satisfied.

sustained using only the supercapacitor but this cost is not expected to be recovered through harvesting in the near future. **Battery-required mission** are those whose energy cost must be totally or partially supplied by the battery.

The classification assigned to the mission does not depend on its profit and demand, but is based on the battery and capacitor energy level, and the node prediction about future energy availability. The same mission can thus be classified in a different way by different nodes.[4] This classification is used by the node to decide whether to bid for a mission or not.

Nodes always accept *free* missions. The energy needed by such missions can be directly provided by the harvesting subsystem without using any energy from the supercapacitor or the battery. This is particularly energy efficient since we avoid energy losses due to supercapacitor inefficiency and leakage. Moreover, in this situation there is no reason for saving energy. Because the capacitor is full, any excess energy harvested from the environment would be wasted if not used.

If the incoming mission is *recoverable* or *capacitor-sustainable*, the node evaluates how profitable the mission is, based on the profit of the mission with respect to the maximum profit and the utility offered by the node with respect to the mission demand. In more detail, a sensor compares the partial profit it can obtain by participating in that mission with the *expected partial profit* $\overline{p}$ of a typical mission, which is computed based on the distribution of the missions profit and demand and the expected utility contribution that a node can offer to a typical mission in its range:

$$\overline{p} = \frac{E[u]}{E[d]} \times \frac{E[p]}{P},$$

where $E[u]$ is the expected utility contribution, $E[d]$ and $E[p]$ [5] are the expected demand and the expected profit of a typical mission, and P is the maximum mission profit.

The partial profit achievable by participating in the incoming mission p^* is defined as:

$$p^* = \frac{u}{d} \times \frac{p}{P} \times w_m, \quad (8)$$

where u is the potential utility contribution that the node can provide to the given mission, d and p are, respectively, the mission demand and profit, P is the maximum mission profit and w_m is the weight associated to the mission's classification (higher for recoverable missions, thus giving a higher chance to those missions to be accepted over capacitor sustained missions). The value p^* is then compared to $\overline{p}$, to have an indication of how profitable the mission is. A node will bid for a given mission only if $p^* \geq \overline{p}$.

Battery required missions are those with the lowest weighting factor. In fact, we want the network to be able to reach a given target lifetime, so we choose to use precious battery energy only to execute missions with higher relative profits. If the mission is *battery-required* the node separately evaluates the energy contribution provided by the supercapacitor and the battery. If the capacitor is not empty, its p^* is computed as in the capacitor-sustainable case and then weighted with the fraction of the requested energy that should be actually provided by the supercapacitor. The same approach is taken to compute the partial profit p^* associated to the battery contribution. However in this latter case we use an additional factor to compute p^*, defined as follows:

$$w_e = \frac{e_a}{e_r},$$

where e_a is the amount of energy available at the node, i.e., its current battery level, and e_r is the amount of energy the node deems to be necessary to reach a given target lifetime.

Let t_e be the *expected occupancy time*, i.e., the fraction of time the node expects to be serving missions in the future. This value is computed based on estimates of the mission arrival rate, the expected mission duration, the probability that a given mission is within the node's sensing range and the probability that a sensors offer is accepted [6]. We denote with τ the remaining target lifetime, i.e, the difference between the initial target network lifetime and the current time. We can then express e_r as: $e_r = \tau \times t_e \times sc$.

The factor w_e is then multiplied for w_m in the computation of p^*. The objective of w_e is to tune the eagerness of sensors to participate in new missions. It forces nodes to be more conservative in accepting missions as the energy they have gets low compared to what is expected to be needed to reach the target lifetime. On the other hand, it also makes the nodes act more aggressively as the target network lifetime approaches.

Since the harvested energy is renewable, and suffers from leakage, there is no point in conserving the energy stored in the capacitor for long periods of time. This is why this criterion, which may appear valid in general, is adopted only to battery-operated missions.

V. Performance Evaluation

In this section we evaluate our sensor-mission assignment scheme EN-MASSE in several different scenarios and compare its performance with the performance of other schemes proposed in the literature.

A. Energy harvesting model

Our C++ simulator uses real-life solar data we collected using Telos B motes [17] interfaced with XOB17-04x3 solar cells [18]. The motes were deployed in a residential area in Rome for a total of 100 non-consecutive days at variable weather conditions and in different locations. A dedicated TinyOS application was developed to track the amount of

[4] The type of the mission can change during its execution (e.g. a recoverable mission can become a sustainable mission if the energy prediction was too optimistic). However, mission classification is used for resource allocation during the bidding phase: if mission classification changes later on this has no effect on node allocation to missions.

[5] All these values are learnt by the nodes based on previous history.

[6] The latter parameter, dubbed γ, is difficult to compute due to a feedback effect: to decide whether a node should bid for a mission we need to know the probability that a bidding is accepted. The approach we have taken has been to experimentally tune the parameter γ so that profit performance of the selected scheme are maximized.

energy generated by the cell every 30 seconds. Due to different weather conditions, seasonal patterns, node position and solar cell orientation, the amount of energy harvested varied significantly over time (as low as 3J per solar cell in case of cloudy and rainy days and up to 220J during sunny days).

EN-MASSE uses an energy prediction model to estimate the amount of energy a node will receive from the ambient source, in order to classify incoming missions, e.g., by predicting the time necessary for the node to recover through harvesting the energy requested by a mission The solar data we recorded, beyond being used as traces for our simulations, allowed us to verify the accuracy of different energy prediction models proposed in the literature, namely the Exponential Weighted Moving Average [11] and the Weather-Conditioned Moving (WCMA) [19] estimation methods. Based on such assessment, we decided to use for EN-MASSE the WCMA energy prediction model, which estimates the energy generation on a typical day based on the energy availability at the same time on the previous days and on the current and past-day weather conditions. This approach in especially effective in case of frequently changing weather conditions, because it provides better accuracy in energy prediction.

B. Simulation setting

In our simulations 500 nodes are randomly and uniformly scattered in a square area of side 400m. The communication range of the nodes is set to 40m. The sensing range is 30m. The node energy model is that of ECO nodes [20], an ultra-compact expandable wireless sensor platform, which has been used in combination with the Ambimax energy harvesting platform [8]. Specifically, ECO node active power consumption is 9mW, while its idle power consumption is 0.006mW. The supercapacitor leakage is modeled as specified in [16], using a set of linear functions approximating the empirical leakage curve. The charging/discharging efficiency of the supercapacitor is set to 95%, i.e., $\eta_c = \eta_d = 0.95$. Missions arrive in the network according to a Poisson arrival process and are assigned to a location randomly and uniformly selected in the deployment area. Mission duration is exponentially distributed with an average mission duration of 1 hour. The mission profit and demand also follow an exponential distribution with average equal to 10 and 2, respectively. We consider a mission satisfaction threshold $T_{sat} = 50\%$, i.e., in order for a mission to be successful, it must receive at least half of its demand from the sensors allocated to it. The utility that a sensor N_i can potentially offer to a mission M_j is defined as a function of the distance D_{ij} between the location l_j of the mission and the position of the node N_i. The network target lifetime T_l has been varied between 30 and 180 days.

In our implementation of EN-MASSE, we set the time threshold within which the energy cost has to be recovered for the mission to be classified as recoverable equal to the mission inter-arrival time times the probability that the new mission falls in the node neighborhood. Thus, whenever a mission is recoverable, bidding for it is not expected to compromise the node capability to serve future missions.

The weighting factors w_m associated with mission types are set equal to 1 for battery required missions, 1.1 for capacitor sustainable missions and 1.2 for recoverable missions.

EN-MASSE was compared to the following energy harvesting unaware mission assignment schemes, proposed in [4] by Johnson et al.: **1) Basic Scheme**: Sensors propose to any mission within their range. **2) Energy Aware Scheme**: This scheme does not use any classification of missions or mechanism to account for energy harvested energy. It simply uses Equation 8 to make bidding decisions, applying a weighting factor w_m set to the ratio between the current residual energy level and the maximum energy level. **3) Energy-Lifetime Aware Scheme**: This is similar to the Energy Aware Scheme, but it also considers the target lifetime of the network when making bidding decisions. It uses Equation 8 to make bidding decisions with a weighting factor equal to the ratio between the time a node can actively sense, given its residual energy level and the target lifetime, and the expected occupancy time of the node.

Because these schemes were not initially designed to use nodes with energy harvesting capabilities, we modified them in order to take both the capacitor and the battery into account when computing the residual and the maximum energy level. The nodes are also able to directly use the energy obtained through harvesting, as well as storing it in the supercapacitor for future use. All the bidding schemes we considered uses the communication protocol described in [3], [4] for the election of the leader and for exchanging information between the mission leader and the nearby nodes; the communication overhead of such a protocol is studied in [3] .

EN-MASSE, the Basic Scheme, the Energy Aware Scheme, and the Energy-Lifetime Aware Scheme have been compared in different scenarios, varying the target lifetime, the supercapacitor and battery size, the mission arrival rate, and the type of sensors embedded in the nodes (thus the energy cost associated to sensing). Here we first discuss results for a specific setting of such parameters taken from a practical case, and then we discuss how the relative performance of the four protocols changes when we explore the whole parameter space.

C. Reference scenario

Our first set of experiments refer to a scenario where node sensing cost is 3 mW (i.e., the power consumption of a temperature sensor like Sensirion SHT11 or a dual-axis accelerometer like ADXL202E). Each node is equipped with two solar cells and with a 25F Maxwell supercapacitor [21], which can store 90J of energy. Using this parameter setting, the battery can sustain 32 hours of continuous monitoring, while the energy contained in the fully-charged capacitor is enough to perform almost 130 minutes of continuous sensing. The target network lifetime is set to four months (shown as a fine vertical line) and the simulations are run for 130 days. The arrival rate of missions in the network is 22 missions per hour. All other parameters are as detailed above.

We evaluate the performance of the different protocols with respect to network profit over time (Fig. 1a), total profit at the

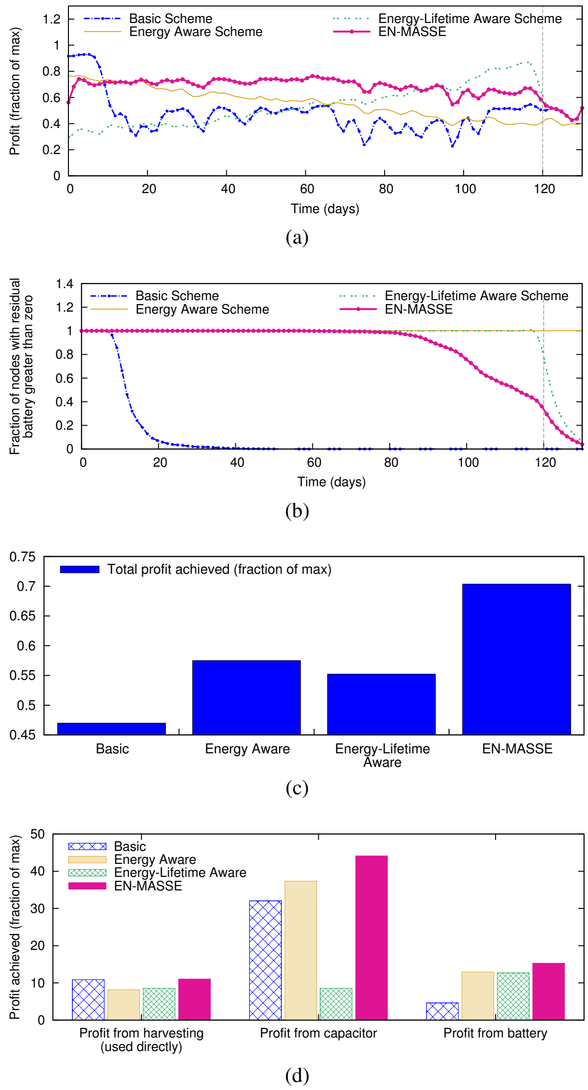

Fig. 1. Simulation results for a sample scenario: (a) fraction of achieved profits per time slot, (b) fraction of nodes with a residual battery greater than zero, (c) total profit achieved at target lifetime (as a fraction of maximum) and (d) total profit achieved per energy source (as a fraction of maximum).

network target lifetime (Fig. 1c) and in terms of fraction of nodes that have energy left in their battery over time (Fig. 1b). The first two metrics allow us to understand which bidding scheme is more profitable and whether a given scheme is able to provide a stable profit over the whole target lifetime or not. The third metric deserves some discussion. Given the energy harvesting capabilities of the nodes, a node with empty battery is not considered dead, as it can execute missions using the energy harvested from the environment. However, a node with no residual battery suffers strongly from fluctuations of the environmental source. It is not able to execute important missions if there are limited energy harvesting opportunities when such missions come in. Therefore having a high percentage of nodes operating only based on the harvesting subsystem may compromise the capability of the network to execute critical missions and degrade the network profit.

We observe that EN-MASSE leads to greatest profit and is able to provide a high stable profit ($70 - 80\%$ of the maximum achievable profit, defined as the profit that could be achieved if all missions were fully served) till the target network lifetime. The fact nodes wisely exploit their resources to be able to provide the needed support to arriving missions is confirmed by Fig. 1b, which shows that no node has depleted its battery in the first 80 days. Then the fraction of nodes with no energy in their battery starts increasing, reaching 60% at the target network lifetime. The remaining nodes, as well as those that can be sporadically charged by the harvesting subsystem, are enough to serve arriving missions with very high profit. This is the desirable behavior for the network. The bidding scheme should not be too aggressive (as is the Basic Scheme) as this would mean that all nodes fast deplete their battery, by thus making the network profit significantly degrade over time. On the other hand, the bidding scheme should not be too conservative, as the energy left in the node batteries at the target network lifetime is wasted. Fig. 1d shows the sources from which EN-MASSE derives its increased profit. Its profit from energy harvesting (both directly and via the capacitor) is larger than the other two methods that are energy aware. Further tests showed that EN-MASSE received more than 40% higher profit than the other energy aware methods from free, recoverable and sustainable missions, indicating its direct consideration of the renewable energy source has large benefits.

The Basic Scheme achieves high profit at the beginning of the network operations. However, since sensors bid for any mission within their range, node batteries start dying rapidly. After 15 days of simulation more than half of the nodes have depleted their batteries and the profit falls below 50% of the maximum profit. After 30 days, less than the 1% of nodes in the network have some residual energy stored in their batteries. From here on, the profit achieved by the scheme shows a high variability, because the capability to serve missions now depends on the amount of harvested energy, which fluctuates as the environmental source does. Overall, as shown in Fig. 1c, which reports the total profit achieved by each scheme over the target network lifetime, the total profit achieved by the Basic Scheme is less than half the maximum and 33% lower than the profit achieved by EN-MASSE.

In the Energy Aware Scheme, having no knowledge of the target network lifetime, each node tries to conserve its resources as long as possible. As can be seen in Fig. 1b, no node has completely depleted its battery after 120 days of simulation. However, to achieve this the scheme has to act conservatively and to ignore many missions. This is why its total profit falls below 60% of the maximum profit.

The Energy-Lifetime Aware Scheme tries to conserve node resources in order to reach the target lifetime. Since it does not exploit information on the harvesting subsystem or on the expected harvested energy, it tends to be overly conservative. This means that, in the first half of the target network lifetime, the support offered to missions (thus the network profit) is quite low. Then the scheme changes its behavior, becoming more aggressive, and increasing its profit. Overall the total profit achieved during the target network lifetime is quite low:

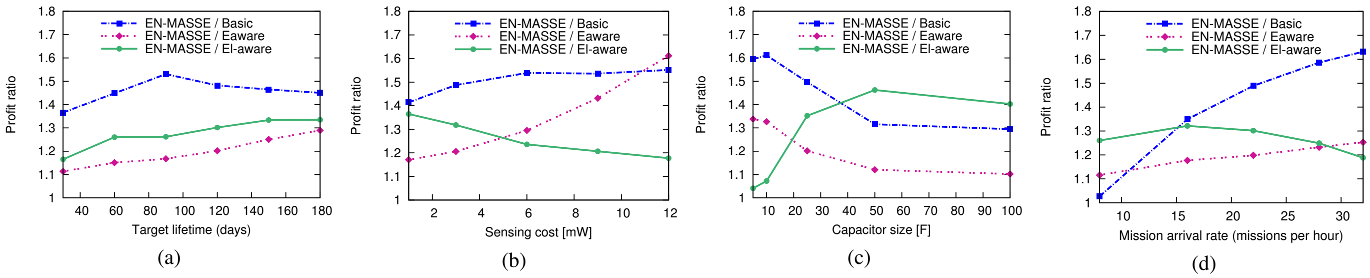

Fig. 2. Simulation results: performance improvement for varying (a) target lifetime, (b) sensing cost, (c) capacitor size and (d) mission arrival rate.

55% the maximum and 25% lower than what achieved by EN-MASSE.

D. *Other scenarios: Impact of parameter variations*

We investigate how varying key parameters affect the performance of EN-MASSE and of the other considered mission assignment schemes. While varying a parameter, the others are set to the following values: the target lifetime is set to 120 days, the mission arrival rate is 22 missions per hour, the sensing cost is 3mW and the capacitor size is 25F; all other parameters are set as specified before. Due to space constraints, for each scenario we do not show the behavior of each scheme over time (as in Section V-C), but only the performance improvement obtained by our EN-MASSE scheme with respect to the other assignment schemes. Fig. 2 shows the ratio between the total profit achieved at target lifetime by EN-MASSE and the profit achieved by the other schemes. All curves have been obtained by averaging results over 10 runs.

1) Target lifetime: As shown in Fig. 2a, the gap in profit between EN-MASSE and the Energy Aware Scheme grows with the target lifetime. The same is true for the Energy-Lifetime Aware scheme. The reason is that for longer network lifetimes a higher percentage of missions are enabled by the harvested energy. This penalizes the Energy Aware and Energy-Lifetime Aware schemes, which do not use exploit information on future harvested energy and on the harvesting subsystem features to make their bidding decisions. Such schemes simply become more and more conservative, loosing in terms of profit, as battery becomes a critical resource, which can satisfy only a small percentage of missions.

The performance ratio of EN-MASSE over the Basic Scheme goes up to 1.5 when the target lifetime is set to 4 months. It then slightly decreases for longer lifetimes. EN-MASSE improved performance reflects the fact that our scheme uses the battery energy to serve missions that have a higher profit than those selected by Basic, and that it selects more profitable missions also when they are enabled by harvested energy. However, the gap in profit between missions satisfied by EN-MASSE and the Basic Scheme is less significant in case of energy harvesting operated missions. Harvested energy must be spent within a limited time frame, making it harder to achieve a high profit in this case. As the target lifetime increases a larger percentage of missions are sustained by energy harvesting, reducing the ratio between the profits achieved by the two schemes.

2) Sensing cost: Fig. 2b shows how the performance of EN-MASSE and the other schemes is impacted by the sensors used and the associated energy consumption. A higher energy consumption for sensing (sensing cost) degrades the performance of the Energy-Aware Scheme. The ratio between EN-MASSE and this scheme can be as high as 1.65 for a sensing cost of 12mW. The reason is that when the sensing cost is high nodes deplete a considerable percentage of their battery energy quickly. In this case, the Energy-Aware Scheme becomes very conservative accepting only missions with very high profit. This would be a good strategy if nodes were equipped only with the battery, as already noticed. When harvested energy is a significant percentage of the overall energy available to the node, being overly conservative in accepting missions is a profit trap, since supercapacitor leakage and finite buffer size reward a fast use of harvested energy. EN-MASSE outperforms the Basic Schemes more significantly as the sensing cost increases. The higher the sensing cost, the higher the toll when making a wrong mission selection.

The performance improvement of EN-MASSE with respect to the Energy-Lifetime Aware Scheme shows a decreasing trend for increasing sensing costs, as this scheme has a finer mechanism to control nodes eagerness to bid for missions. Despite of that, EN-MASSE still outperforms the Energy-Lifetime Aware Scheme, achieving a profit ratios between 1.35 (smaller sensing cost) and 1.2 (higher ones).

3) Capacitor size: As shown in Fig. 2c, the performance improvement of EN-MASSE with respect to the Basic and Energy Aware schemes has a decreasing trend for increasing capacitor size. The Basic Scheme takes advantage of a bigger capacitor. Basic usually suffers from a dumb management of the supercapacitor energy buffer; when such buffer is larger this effect is less important. The Energy Aware Scheme shows a similar trend because, when making proposal decisions, it takes into account the ratio between the current residual energy (battery plus capacitor) and the maximum energy. The bigger the capacitor, the more aggressive the scheme is when harvested energy is available.

The performance improvement of EN-MASSE with respect to the Energy-Lifetime Aware Scheme instead increases for bigger capacitor sizes, demonstrating that our scheme better exploits recharge opportunities. We also observe that the relative performance of the assignment schemes does not show a significant variation when using a 100F capacitor instead of

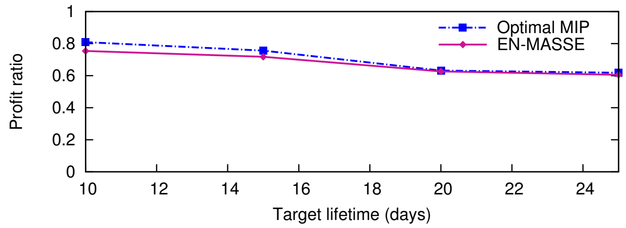

Fig. 3. Performance of our EN-MASSE with respect to the upper bound provided by the optimal MIP solution.

a 50F capacitor, because of its higher leakage.

4) Mission arrival rate: Fig. 2d shows how the performance of EN-MASSE varies with respect to the other schemes for different mission arrival rates. The performance improvement over the Energy Aware Scheme increases for higher mission arrival rates up to a ratio of 1.3. In fact, increasing the mission arrival rate has an effect similar to considering higher sensing cost, as it leads to heavier network workload. Thus, once again, the Energy-Aware Scheme becomes conservative and accepts only missions with high profit. Since nodes using the Basic Scheme propose to any mission within their range, while EN-MASSE selects missions with higher profit, the gap in profit between the two schemes grows up to a ratio of 1.65 as the mission arrival rate increases and more missions arrive in the network. Finally, the performance improvement over the Energy-Lifetime Aware Scheme slightly decreases for increasing mission arrival rates, but it remains between a ratio of 1.35 and 1.2.

E. Comparison with the optimal MIP solution

We compare our solution to the optimal mixed integer programming (MIP) solution, obtained by solving a formulation of the problem where the mission success threshold is set to 0 and energy buffers are ideal (no leakage, infinite buffer, charge and discharge efficiency set to 1). Solving this model provides an upper bound on the solution to our original formulation.

We consider a network with 25 nodes and a mission arrival rate of 4 missions per hour. All other simulation parameters (including the battery size) have been scaled accordingly.

Fig. 3 shows the performance of EN-MASSE with respect to the upper bound provided by the optimal MIP solution for different target lifetimes. The y-axis shows the total profit obtained at target lifetime as a fraction of the maximum profit. The gap between the two solutions is around 5% of the maximum profit and it decreases for longer target lifetimes.

VI. Conclusions

In this paper we have presented an analytical model and a distributed solution, EN-MASSE, for sensor-mission assignment in WSNs with energy harvesting. Our distributed scheme, EN-MASSE, is shown to perform very closely to the optimum provided by the analytical formulation, and to outperform other mission assignment solutions. In particular, by comparing mission assignment schemes in several different scenarios we have demonstrated that traditional assignment algorithms cannot harness the full potential provided by the harvesting technology, which is instead taken into account efficiently by our proposed scheme.

Acknowledgments

Dora Spenza is a recipient of the Google Europe Fellowship in Wireless Networking, and this research is supported in part by this Google Fellowship. This research was also sponsored in part by the US National Science Foundation under Grant CNS-0916171 and by the FP7 project GENESI (GrEen sensor NEtworks for Structural monItoring).

References

[1] H. Rowaihy, S. Eswaran, M. Johnson, D. Verma, A. Bar-noy, and T. Brown, "A survey of sensor selection schemes in wireless sensor networks," in *SPIE Defense and Security Symposium Conference*, 2007.

[2] H. Rowaihy, M. P. Johnson, D. Pizzocaro, A. Bar-Noy, L. M. Kaplan, T. F. L. Porta, and A. D. Preece, "Detection and localization sensor assignment with exact and fuzzy locations," in *DCOSS*, 2009, pp. 28–43.

[3] H. Rowaihy, M. P. Johnson, A. Bar-Noy, T. Brown, and T. F. L. Porta, "Assigning sensors to competing missions," in *GLOBECOM*, 2008, pp. 44–49.

[4] M. P. Johnson, H. Rowaihy, D. Pizzocaro, A. Bar-Noy, S. Chalmers, T. F. L. Porta, and A. D. Preece, "Frugal sensor assignment," in *DCOSS*, 2008, pp. 219–236.

[5] A. Bar-Noy, T. Brown, M. P. Johnson, T. F. L. Porta, O. Liu, and H. Rowaihy, "Assigning sensors to missions with demands," in *ALGOSENSORS*, 2007, pp. 114–125.

[6] T. Le, T. J. Norman, and W. W. Vasconcelos, "Agent-based sensor-mission assignment for tasks sharing assets," in *Third International Workshop on Agent Technology for Sensor Networks*, 2009.

[7] X. Jiang, J. Polastre, and D. Culler, "Perpetual environmentally powered sensor networks," *IPSN 2005*, pp. 463–468, April 2005.

[8] C. Park and P. H. Chou, "Ambimax: Autonomous energy harvesting platform for multi-supply wireless sensor nodes," in *SECON '06*, vol. 1, Sept. 2006, pp. 168–177.

[9] V. Raghunathan, A. Kansal, J. Hsu, J. Friedman, and M. Srivastava, "Design considerations for solar energy harvesting wireless embedded systems," in *IPSN '05*, 2005, p. 64.

[10] F. Simjee and P. H. Chou, "Everlast: long-life, supercapacitor-operated wireless sensor node," in *ISLPED '06*, 2006, pp. 197–202.

[11] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava, "Power management in energy harvesting sensor networks," *ACM Trans. Embed. Comput. Syst.*, vol. 6, no. 4, p. 32, 2007.

[12] C. Vigorito, D. Ganesan, and A. Barto, "Adaptive control of duty cycling in energy-harvesting wireless sensor networks," *SECON '07*, pp. 21–30, June 2007.

[13] J. Recas Piorno, C. Bergonzini, D. Atienza Alonso, and T. S. Rosing, "HOLLOWS: A Power-Aware Task Scheduler for Energy Harvesting Sensor Nodes," *Journal of Intelligent Material Systems and Structures*, vol. 21, no. 12, pp. 1300–1332, 2010.

[14] C. Moser, L. Thiele, D. Brunelli, and L. Benini, "Adaptive power management in energy harvesting systems," in *DATE '07*, 2007, pp. 773–778.

[15] J. B. Steck and T. S. Rosing, "Adapting task utility in externally triggered energy harvesting wireless sensing systems," in *INSS'09*, 2009, pp. 16–23.

[16] T. Zhu, Z. Zhong, Y. Gu, T. He, and Z.-L. Zhang, "Leakage-aware energy synchronization for wireless sensor networks," in *MobiSys 09*, 2009, pp. 319–332.

[17] Crossbow Technology, "TelosB mote platform datasheet," document Part Number: 6020-0094-01 Rev B.

[18] IXYS Corporation, "XOB17 IXOLAR High Efficiency Solar Bits technical information," Jun 2009.

[19] J. Recas Piorno, C. Bergonzini, D. Atienza, and T. Simunic Rosing, "Prediction and management in energy harvested wireless sensor nodes," in *Wireless Vitae 2009*, 2009.

[20] C. Park, J. Liu, and P. H. Chou, "Eco: an ultra-compact low-power wireless sensor node for real-time motion monitoring," in *IPSN '05*, 2005, p. 54.

[21] Maxwell Technologies, "Datasheet HC power series ultracapacitors," August 2009.