

Sensor-mission assignment in rechargeable wireless sensor networks

THOMAS LA PORTA, Pennsylvania State University
CHIARA PETRIOLI, Sapienza University of Rome
CYNTHIA PHILLIPS, Sandia National Laboratories
DORA SPENZA, Sapienza University of Rome

Sensor mission assignment involves 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 and are powered in part by uncontrollable environmental sources, 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, specifically tailored for energy-harvesting mission-centric WSNs. To assess the performance of our proposed solution we have interfaced TelosB nodes with solar cells and performed extensive experiments to derive models and traces of solar energy acquisition. We use such real-life traces in our simulations. A comparative performance evaluation between EN-MASSE and other schemes previously proposed in the literature has shown that our solution significantly outperforms existing energy-harvesting-unaware mission assignment schemes. Moreover, using our analytical model as a benchmark, we also show that the profit earned by EN-MASSE is close to the optimum. Finally, we have implemented our proposed solution in TinyOS and experimentally validated its performance, showing the effectiveness of our approach.

Categories and Subject Descriptors: C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design

General Terms: Design, Performance, Measurement, Experimentation

Additional Key Words and Phrases: Wireless sensor networks, sensor-mission assignment, energy harvesting, supercapacitor, solar, rechargeable, task allocation, harvesting aware, optimization, MIP

ACM Reference Format:

Thomas La Porta, Chiara Petrioli, Cynthia Phillips and Dora Spenza, 2013. Sensor-mission assignment in rechargeable wireless sensor networks. *ACM Trans. Sensor Netw.* V, N, Article A (January YYYY), 40 pages. DOI: <http://dx.doi.org/10.1145/0000000.0000000>

A preliminary version of this article was presented at the 8th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON 2011) and appears in the conference proceedings [La Porta et al. 2011].

Dora Spenza is a recipient of the Google Europe Fellowship in Wireless Networking. This research is supported in part by this Google Fellowship. This work 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). Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Author's addresses: T. La Porta, Computer Science and Engineering Department, Pennsylvania State University; C. Petrioli and D. Spenza, Computer Science Department, Sapienza University of Rome; C. Phillips, Sandia National Laboratories, Albuquerque, NM.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 1550-4859/YYYY/01-ARTA \$15.00
DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

Field and environmental monitoring are common applications of wireless sensor networks (WSNs). In a typical scenario, the nodes of a WSN are placed in a region for collecting 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 completed successfully. There is a tradeoff between accuracy, the number of tasks 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 [Bartolini et al. 2012; Rowaihy et al. 2007; Shih et al. 2006] 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.

Many applications, however, may require the network to achieve multiple, simultaneous missions. By *mission*, we refer to a sensing task whose primary goal is the collection of information, to which one or multiple sensors may contribute. In general, sensor nodes, together with other information sources, may be part of a collection of intelligence, surveillance, and reconnaissance (ISR) assets that should be allocated to support the objectives of a mission [Preece et al. 2008]. In application scenarios such as battlefield monitoring and emergency disaster response, the informational demands placed on ISR resources typically exceeds their supply in terms of inventory [Gomez et al. 2008], which results in simultaneous missions competing for the sensing resources of the nodes. For example, consider an intruder detection application in which nodes are equipped with pan-tilt-zoom cameras that can be rotated to point in a direction of interest: the same camera can not be used by two different missions that require monitoring two opposite regions of a field [Pizzocaro et al. 2008]. A similar competition may occur in a network deployed to perform concurrent localization and event detection tasks using directional acoustic sensors, which can be assigned to only one mission at a time [Rowaihy et al. 2009]. In general, for these mission-centric wireless sensor networks, a sensor selection scheme is no longer sufficient, because it is necessary not only to decide which nodes in the network are active, but also to assign each active sensor to the mission it may serve best. This is a non trivial task, because as a given node may offer support to different missions with different levels of accuracy and fit (*utility*). Meanwhile, missions 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 only if a sufficiently good set of node is assigned to it, depending on its demand and on the quality of contribution that the assigned sensors can provide. The goal of a sensor-to-mission assignment algorithm is to assign available nodes to appropriate missions, maximizing the profit received by the network for mission execution.

Different constraints on the assignment decisions lead to different versions of this problem, depending on whether missions arrive over time [Rowaihy et al. 2009; Rowaihy et al. 2008; Johnson et al. 2010] (dynamic case) or they are all available when network operations start [Bar-Noy et al. 2008; Rowaihy et al. 2008; Johnson et al. 2010] (static case). Other variants concern whether the network has a predefined target operational lifetime [Johnson et al. 2010] or it must just last as long as possible [Rowaihy et al. 2008]. Some prior solutions proposed in previous works 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 with traditional batteries, mount *energy-harvesting* devices that opportunistically draw new energy from the environment [Basagni et al. 2013].

Unlike traditional motes, nodes with energy harvesting capabilities alternate between periods in which energy must be sparingly used, and situations in which there may be an excess of energy available, which would be wasted unless used in the short term [Bianchi et al. 2013]. 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. These solutions should also consider that energy availability is time-dependent, and that energy storage is limited in size and time (due to self-discharge). So energy usage should be carefully planned to minimize energy waste.

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, serving missions that arrive dynamically in the network (*dynamic assignment with a time horizon*). In our study, we consider sensor nodes with 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 make the following contributions:

- We model the problem of optimum mission assignment in energy-harvesting mission-centric WSNs. By addressing scalability issues of previous approaches, we provide a MIP formulation that can be solved for real-life-sized instances, consisting of hundreds of nodes and with target lifetimes of several months. At the same time, our MIP model captures the details of the behavior of a typical energy harvesting subsystem. The mathematical model we propose provides an upper bound to the maximum profit achievable by the networks, serving as a benchmark for more realistic, distributed protocols.
- 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 considers not only the nodal residual energy, but also the energy it expects to harvest in the future (using an energy prediction model) and the expected profit and demand of missions to come.
- 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 several months under variable weather conditions and in different locations. We also validate our approach by using four additional solar datasets obtained from the US Climate Reference Network [USCRN 2011].
- We perform a comparative performance evaluation of EN-MASSE and of the schemes proposed by Johnson et al. in [Johnson et al. 2010]. We evaluate the impact of critical parameters (such as the target lifetime, the energy consumption of sensors embedded in the nodes, the supercapacitor size, and the mission arrival rate) on the performance of the different schemes in twenty distinct scenarios. Furthermore, we show that the profit earned by EN-MASSE is remarkably close to the upper bound on the optimum obtained through our MIP model.
- Finally, we implemented our proposed solution in TinyOS and experimentally validate its performance in a real-life testbed of solar-powered Telos B motes equipped with cameras.

The remainder of this paper is organized as follows. In Section 2 we formalize the mission assignment problem for energy-harvesting WSNs, and provide a MIP formulation

that maximizes network profit. We present EN-MASSE in Section 3, and compare its performance to that of previous solutions and to the upper bound derived by the MIP model in Section 4. In Section 5 the experimental validation of EN-MASSE is presented, reporting results obtained from a testbed of Telos B motes interfaced with directional video sensors. We discuss related work in Section 6. Finally, we present our conclusions in Section 7.

2. PROBLEM FORMULATION

In this section, we formulate the mission assignment problem for energy-harvesting WSNs. We begin by describing the architecture of a typical node with energy harvesting capability and the general application scenarios for our model. We then discuss our analytical model, and provide a MIP formulation for sensor-mission assignment in energy-harvesting WSNs. Finally, we briefly discuss hardness results for the considered problem.

2.1. Node architecture

Each node in the network is equipped with an energy harvesting subsystem, which includes one or more photo-voltaic panels¹ and a supercapacitor that acts as an energy buffer for the node. We model several realistic characteristics of a typical supercapacitor, including the fact that it has a finite size C^{Max} , that it suffers from leakage and self-discharge, and that it has a charging efficiency $\eta_c^C < 1$ and a discharging efficiency $\eta_d^C < 1$. Nodes are also equipped with a non-rechargeable primary battery of capacity B^{max} , whose purpose is to guarantee a minimum lifetime in case no or little energy can be drawn from the environment. Similar to the supercapacitor, the battery also has a discharging efficiency $\eta_d^B < 1$.

The hybrid system of solar harvester, supercapacitor, and primary battery provides the energy to power the node and to support the sensing activity required for mission execution.

2.2. Application scenarios

As briefly introduced in Section 1, we wish to provide a formulation of the sensor-mission assignment problem that captures several aspects of real-life systems, while providing enough flexibility to be applicable to a number of different scenarios.

Our model allows us to define:

- (1) an application-specific measure of how useful the data generated by a particular node are for performing a given task (utility);
- (2) the maximum level of utility a task may require from the sensor network (mission demand);
- (3) the priority and importance of each task (mission profit);

The utility function indicates the quality of the information that a node can provide to a mission, giving an estimation of how well it would fit the purpose of the mission. Such a concept can be employed to define several application-specific requirements of practical scenarios [Bian et al. 2006]. For example, a microphone may be more useful to the application if it is closer to an area of interest. Thus, the utility it can provide may be defined as a function of its distance from the observed area. Similarly, a given camera node may be totally useless to a surveillance application if its video sensor is not oriented in a given direction. In this case, the node utility would be zero, unless its orientation matches the application requirement.

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

The demand of a mission indicates the amount of sensing resource capabilities it requires. Intuitively, it represents the “difficulty” of a mission, as tasks with higher demands require greater amount of resources to be executed. The demand can be used to simply indicate how many sensors should be assigned to a particular mission, but also to evaluate the overall quality of contribution that the nodes assigned to a task can provide to it. For instance, a mission to capture a 3D representation of a particular target will fail if only one camera is assigned to it. Moreover, the accuracy of the resulting image will depend not only on how many nodes are assigned to the mission, but also on their positions and orientations. In other scenarios, instead, it may be desirable to have as many nodes as possible monitoring the same target to increase the redundancy of the measurements and, thus, the robustness of the network to unpredictable events [Chen et al. 2012]. In general, the interpretation of the demand strictly depends on the specific metric used to calculate utility values.

Finally, the profit of a task is an application-specific quantity that defines its level of importance. This quantity may be used to discriminate between routine, low-profit missions, and critical, high-profit missions. For example, high-profit missions can be generated when measured values are above an alarm threshold, or upon occurrence of specific events. If the network is shared among multiple users, priority may also be assigned to missions generated by users based on their role.

It is worth noting that profit and demand of a mission are not necessarily correlated. A low-priority monitoring mission may be expensive in terms of sensing demand. At the same time, an urgent mission that requires information about a limited area in which an alarm has been triggered may be low-demand, but, being critical for the application, have very high profit.

2.3. Analytical model

We formulate the mission assignment problem for wireless sensor networks with energy harvesting capabilities as follows. The network consists of a set of energy-harvesting-endowed sensor nodes N_1, \dots, 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, \dots, M_m be the set of missions that the network is asked to perform.

A mission M_j is a tuple $(p_j, d_j, g_j, ts_j, te_j)$ where:

- p_j . is the profit of the mission per unit time, indicating both its importance and the reward achieved by the network for its accomplishment;
- d_j . is the mission demand, indicating the amount of sensing resources it needs;
- g_j . is the geographic location of the mission in the field;
- ts_j . is the time at which the mission arrives in the network;
- te_j . is the time at which the mission terminates.

Missions usually last for multiple units of time, i.e., $te_j > ts_j$. During mission execution, p_j, d_j, g_j remains constant over time.

A mission is executed only if a sufficiently good set of node is assigned to it, depending on its demand and on the quality of contribution that the assigned sensors can provide. We define e_{ij} as the utility received by mission M_j if node N_i is assigned to it. This utility is zero if the node cannot contribute to the mission, which happens, for instance, if the node is not close enough to the mission location or if the orientation of the directional sensor is not useful to the mission.

While missions can be performed by multiple sensors simultaneously, we assume that a node, being equipped with directional sensors, can be assigned to at most one mission at a time.

Let u_{jt} denote the total utility received by a mission M_j at time t . We express the total utility u_{jt} as the sum of the utilities provided by the sensors assigned to the

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




Fig. 1. The profit achieved for mission execution is a function of mission satisfaction level

mission M_j at time t . That is, $u_{jt} = \sum x_{ijt}e_{ij}$, where $x_{ijt} = 1$ if and only if node i is serving Mission j at time t . In this model, utility from each node is independent of the other nodes assigned to a sensor. Although this weighted linear model does not capture the 3D representation example of Section 2.2, it does capture the inability of specific sensors to serve a mission as described in the same section. Our model represents utility models more complex than simple counting, as well as other scenario-specific utility functions defined by the user.

The total utility that a mission M_j requires is expressed by its demand d_j . If the total utility u_{jt} received by a mission is lower than its requested demand d_j , we say that the mission M_j is partially satisfied (at time t) and we indicate its satisfaction level with $y_{jt} = u_{jt}/d_j$ (in the range $[0, 1]$).

Profits are received by the network for mission accomplishment, 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. We use a single satisfaction threshold for all missions, but the formulation would be essentially the same if the threshold varies by mission.

The profit achieved for executing mission M_j at time t depends on the satisfaction level y_{jt} of the mission at time t , as shown in Figure 1. The profit achieved by the network for mission execution is: zero if the minimum satisfaction threshold T_{sat} is not met; a fraction of the mission profit if the satisfaction threshold is met, but the mission is not fully satisfied; equal to the mission profit, p_j , if the mission is fully satisfied.

The total profit received for executing mission M_j is the sum of the profits earned over the entire mission lifetime, i.e., $p_j = \sum_{t=t_{s_j}}^{t_{e_j}} p_{jt}(y_{jt})$.

Ideally, we seek an assignment of sensors to missions that always satisfies each mission's demand at least to the given satisfaction threshold. 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.

2.4. Sensor-mission assignment problem: MIP formulation

In this section, we describe a MIP formulation for the sensor-mission assignment problem in energy-harvesting WSNs, whose solution provides an upper bound on the maximum achievable profit over a given target lifetime.

Previous works have proposed MIP formulation for sensor-mission assignment in traditional WSNs. For example, Johnson et al. considered battery-operated nodes that lose energy monotonically [Johnson et al. 2010]. When nodes can harvest energy, how-




Fig. 2. Power harvested by the node N_i during time epoch t and energy consumption required by the sensing activity.

ever, the energy available to each node fluctuates over time. This characteristic adds significant complexity to the MIP problem, because it is necessary to update the energy level of the nodes with higher frequency to accurately represent the current energy availability. The model should also capture energy buffer non-idealities, such as limited capacity and charging/discharging efficiency, to avoid grossly overestimating the energy available to the nodes.

One possible solution to this problem is to divide time into discrete time epochs of small size and update the energy level of the nodes periodically [La Porta et al. 2011]. Since the environmental energy strongly varies over time, such updates should be performed within seconds or minutes, with a tradeoff between an accurate representation of the energy variations and the model scalability. However, we found such an approach difficult to scale, especially when considering real-life instances with hundreds of nodes and target lifetimes of weeks or months.

In this work we propose a more scalable approach to formalize the sensor-mission assignment in energy-harvesting WSNs. Our idea is to divide time into discrete time epochs of relatively large size, and to compute off-line the variations in the energy level of the nodes within each of such epochs. More specifically, we define the set of time epochs in which the energy level of the node N_i is updated, τ_i^c (where c stands for “check”), as the union of all times in which a mission starts or ends within its sensing range. In this way, when a new mission arrives and a decision has to be made about assigning a node to it or not, we ensure that the energy level of the node is up to date. Similarly, at the end of a mission the energy level of each node that was assigned to it is updated, considering the energy spent for mission execution. Figure 3 (page 11) shows an example of τ_i^c : missions M_i, M_j arrive within the sensing range of node N_i at different times. The energy level of N_i is updated at both their start and end times.

The variations in the energy level of the nodes are not explicitly calculated in the model; instead, they are pre-computed off-line through simulations that account for both fluctuations in the environmental energy availability and non-ideal behaviors of the energy storage. These pre-computations are necessary because the energy available to the node fluctuates over time, according to the variability of the energy source. For this reason, considering a simple indicator, such as the average power, of the energy harvested during each time epoch is not enough to correctly represents the real energy availability. For instance, Figure 2 shows the power harvested by node N_v dur-

ing a particular time epoch t spanning between 1:30PM and 2:40PM. During this period, the mean harvested power is around 14.20 mW, while the power consumption of the node is set to 12 mW, assuming N_i is executing some mission M_j . Despite the fact that the mean harvested power is higher than the power consumption of N_i , at the beginning of time epoch t the harvested power is not enough to directly power the node, thus requiring some energy from an external buffer to be supplied. To correctly model such situations, the MIP model should assign the node N_i to the mission M_j only if N_i has enough energy available in its battery and supercapacitor at the beginning of time epoch t .

Energy pre-computations allow to identify such critical points in terms of energy availability and to make correct decisions for mission assignment, even when considering large time epochs during which the harvested energy may vary significantly. Moreover, pre-computations are used to simulate how the energy level of the supercapacitor will vary during each time epoch, allowing to capture features of realistic energy buffers, such as limited capacity and charging/discharging efficiency. More details about this approach are explained in Section 2.4.2.

Compared to fine-granularity updates of the energy level of the node, our solution provides a substantial reduction of the number of time epochs considered in the model, thus significantly improving the MIP problem scalability.

2.4.1. MIP formulation. We model the sensor-mission assignment problem in energy-harvesting WSNs as reported in the mathematical program (MP) 1. Please refer to Table I for notations.

We seek an assignment of nodes to missions that maximizes the total profit obtained by the network (1). There are six sets of variables:

- y_{jt} . is the satisfaction level of the mission M_j at time t ;
- z_{jt} . indicates whether y_{ij} is 0 or it is in its continuous region;
- x_{ijt} . indicates if the node N_i starts serving the mission M_j at time t ;
- C_{it} . is the energy level of the supercapacitor of the node N_i at time t ;
- B_{it} . is the energy level of the battery of the node N_i at time t ;
- D_{it} . is the energy drained from the battery of the node N_i during time epoch t .

The x_{ijt} variables are the only decision variables. The others are determined by the x_{ijt} variables, the other constraints, and the objective function.

The satisfaction level y_{jt} of the mission M_j at time t depends on the utility received by the mission in respect to its demand (2). The variable y_{jt} is effectively a semi-continuous variable, which either must take a value between T_{sat} and 1, or it must be zero (constraints (12) enforces the upper and lower bounds). We implement this using the binary helper variables z_{ij} . If $z_{ij} = 0$, then $y_{ij} = 0$ by constraints (9). Constraints (10) ensure that z_{ij} is not allowed to be 1, and hence y_{ij} is not allowed to be non-zero, until $y_{ij} \geq T_{SAT}$. We consider the case in which preemption between missions is not allowed: a node may be assigned to at most one mission at a time and, once assigned to it, it runs the mission until completion (3). A node can only start executing a mission either when it arrives in the network or when another mission that the node was executing ends. Thus, nodes are forbidden to start a mission at an arbitrary time within its duration. Problem constraints (4) implement this rule, as they do not allow a node to start serving a mission that already started, unless another mission that the node was serving just terminated.

Constraints (5) update the energy level of the capacitor after each time epoch, both in the case the node is executing a mission, or it is idle. When a node starts executing a mission, it must have enough energy stored in its capacitor and in its battery to

$$\max \sum_{M_j} \sum_{t \in Tm_j} L_j(t) * p_j * y_{jt} \quad (1)$$

$$\text{s.t.} \quad \sum_{\substack{t' \in Tm_j: \\ t' \leq t}} \sum_{\substack{N_i: \\ j \in aM_i(t')}} x_{ijt'} * e_{ij} \geq d_j * y_{jt}, \quad \forall M_j, t \in Tm_j \quad (2)$$

$$\sum_{j \in aM_i(t)} \sum_{\substack{t' \in \tau_i: \\ ts_j \leq t' \leq t}} x_{ijt'} \leq 1, \quad \forall N_i, t \in \tau_i \quad (3)$$

$$x_{ijt} \leq \sum_{\substack{j' \in aM_i(t): \\ te_{j'}=t}} \sum_{\substack{t' \in \tau_i: \\ ts_{j'} \leq t' < t}} x_{ij't'}, \quad \forall N_i, t \in \tau_i, j \in aM_i(t) : ts_j < t \quad (4)$$

$$C_{i\hat{t}} \leq C_{it} + D_{it} + \Delta_{it}^I + (\Delta_{it}^I - \Delta_{it}) \sum_{j \in aM_i(t)} \sum_{\substack{t' \in \tau_i: \\ ts_j \leq t' \leq t}} x_{ijt'}, \quad \forall N_i, t \in \tau_i^c \quad (5)$$

$$B_{i\hat{t}} \leq B_{it} - D_{it}, \quad \forall N_i, t \in \tau_i^c \quad (6)$$

$$D_{it} \leq B_{it}, \quad \forall N_i, t \in \tau_i^c \quad (7)$$

$$C_{it} + B_{it} \geq \mu_{ijt} x_{ijt}, \quad \forall N_i, t \in \tau_i, j \in aM_i(t) \quad (8)$$

$$y_{jt} \leq z_{jt} \quad \forall M_j, t \in Tm_j \quad (9)$$

$$z_{jt} \leq 1 + y_{jt} - T_{sat} \quad \forall M_j, t \in Tm_j \quad (10)$$

$$x_{ijt} \in \{0, 1\}, \quad \forall N_i, t \in \tau_i, j \in aM_i(t) \quad (11)$$

$$0 \leq y_{jt} \leq 1 \quad \forall M_j, t \in Tm_j \quad (12)$$

$$z_{jt} \in \{0, 1\} \quad \forall M_j, t \in Tm_j \quad (13)$$

$$0 \leq C_{it} \leq C_{max} \quad \forall N_i, t \in \tau_i^c \quad (14)$$

$$0 \leq B_{it} \leq \eta_d^B B_{max} \quad \forall N_i, t \in \tau_i^c \quad (15)$$

$$0 \leq D_{jt} \leq \eta_d^B B_{max} \quad \forall N_i, t \in \tau_i^c \quad (16)$$

MP 1: Sensor-mission assignment problem in energy-harvesting WSNs

run it until completion (8). Problem constraints (5) and (8) are explained in details in Section 2.4.2.

Constraints (6) update the level of the primary battery when the node drains some energy from it. For every node and every time epoch, the energy drained from the battery must not exceed the available energy (7). Finally, nodes can not be fractionally assigned to a mission (11) and both the supercapacitor and the battery have a finite size (see variable definitions (14) and (15)). Since the primary battery has a discharge efficiency strictly lower than one, its maximum capacity is scaled by a factor equal to η_d^B . As such a battery is non-rechargeable, no charging efficiency is modeled for it. The

Table I. Table of notations.

Symbol	Explanation
Variables	
y_{jt}	satisfaction level of mission M_j at time t
z_{jt}	(binary) indicates whether y_{ij} is 0 or in its continuous region
x_{ijt}	(binary) indicates if node N_i starts serving mission M_j at time t
C_{it}	supercapacitor energy level of node N_i at time t
B_{it}	battery energy level of node N_i at time t
D_{it}	energy drained from battery of node N_i during time epoch t
Parameters	
p_j	profit of the mission M_j
d_j	demand of the mission M_j
g_j	geographic location of the mission M_j in the field
ts_j	time mission M_j arrives in the network
te_j	time mission M_j terminates
e_{ij}	utility received by mission M_j if node N_i is assigned to it
u_{jt}	total utility received by the mission M_j at time t
T_{sat}	minimum satisfaction threshold; if $y_{jt} < T_{sat}$ no profit is awarded for mission execution
C_{max}	supercapacitor size
B_{max}	battery size
η_c^C	supercapacitor charging efficiency, $\eta_c^C < 1$
η_d^C	supercapacitor discharging efficiency, $\eta_d^C < 1$
η_d^B	battery discharging efficiency, $\eta_d^B < 1$
τ_i	set of time epochs in which node N_i can start executing a mission
τ_i^c	set of time epochs in which the energy level of the supercapacitor of the node N_i is updated, i.e., $\tau_i^c = \bigcup_{j \in aM_i} \{ts_j \cup te_j\}$
$aM_i(t)$	set of missions node N_i may serve at time epoch t , i.e. $aM_i(t) = \{M_j : e_{ij} > 0 \wedge ts_j \leq t \leq te_j\}$
aM_i	set of missions node N_i can serve, i.e. $aM_i = \bigcup_{t \in \tau_i} aM_i(t)$
Tm_j	set of time epochs in which a node may start executing mission M_j i.e. $Tm_j = \bigcup_{N_i: e_{ij} > 0} \{t \in \tau_i : ts_j \leq t \leq te_j\}$
\hat{t}	next time epoch after t , i.e., $S = (t_1, \dots, t, t_{i+1}, \dots, t_s)$, $\hat{t} = t_{i+1}$
$L_j(t)$	length of the time epoch t in the set Tm_j , i.e., $L_j(t) = \min(\hat{t}, te_j) - t$
Δ_{it}^c	variation in the energy level of the supercapacitor of the node N_i during time epoch t , $t \in \tau_i^c$, assuming the node is executing a mission
Δ_{it}^l	variation in the energy level of the supercapacitor of the node N_i during time epoch t , $t \in \tau_i^c$, assuming the node is not executing a mission
μ_{it}	minimum energy level reached by the supercapacitor of the node N_i during time epoch t , $t \in \tau_i^c$
μ_{ijt}	minimum energy level reached by the supercapacitor of the node N_i if N_i starts executing M_j during time epoch t , $t \in \tau_i^c$

charging and discharging efficiencies of the supercapacitor are not explicitly modeled in the MIP problem, as we account for them during energy pre-computations.

2.4.2. Supercapacitor non-idealities. Based on current energy-harvesting node prototypes [Magno et al. 2012; Kansal et al. 2007], we model energy harvesting subsystems with the following real-life characteristics:

- (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 self-discharge [Kansal et al. 2007].
- (2) If the amount of energy harvested is greater than the current energy consumption, some energy is directly used to sustain the node's operation, while excess energy is stored, if possible, in the supercapacitor for later use.

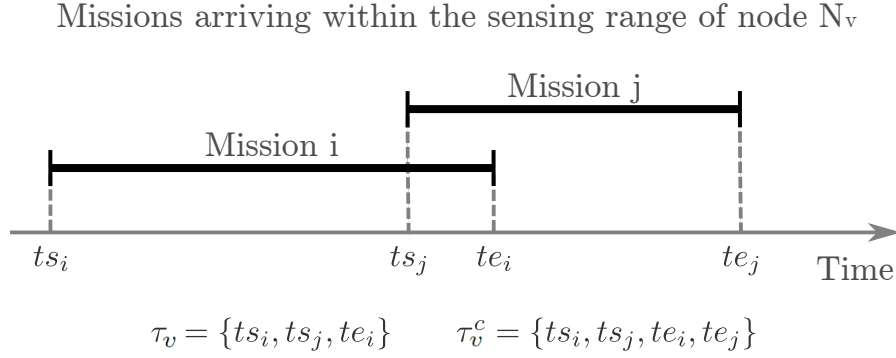


Fig. 3. Missions M_i and M_j appear at different times within the sensing range of node N_v . N_v can start executing mission M_i at time ts_i . It can start executing missions M_j either at time ts_j or te_i (if it was executing mission M_i before M_j arrived). Thus, $\tau_v = \{ts_i, ts_j, te_i\}$. The energy level of node N_v is updated before each mission's start and end times; $\tau_v^c = \{ts_i, ts_j, te_i, te_j\}$.

Because of the finite size of the buffer, some energy may be lost if there is not enough space left in the supercapacitor to store it. The definition of the variables C_{it} in (14) ensures that the supercapacitor is never charged beyond its maximum size.

Furthermore, because both the charging and discharging efficiency of the buffer are strictly less than one, only a fraction $\eta_c^C \cdot \eta_d^C$ of the excess energy is available after storing (and retrieving) it. We account for those non-idealities in our energy pre-computations, by allowing the node to consume the harvested energy directly, when possible, and by storing in the supercapacitor only a fraction $\eta_c^C \cdot \eta_d^C$ of the excess energy.

For each node N_i and for each time epoch $t, t \in \tau_i^c$, we simulate both N_i executing a mission and being idle. In fact, during energy pre-computations, it is not possible to know if, according to the solution of the optimization problem, node N_i will be assigned to a mission M_j during time epoch t .

In order to track the current energy level of the node, and to ensure that the energy required to execute a mission is available before assigning the node to the mission, we define the following parameters:

- Δ_{it} . is the change in the energy level of the supercapacitor of node N_i during time epoch $t, t \in \tau_i^c$, assuming the node is executing a mission;
- Δ_{it}^I . is the change in the energy level of the supercapacitor of node N_i during time epoch $t, t \in \tau_i^c$, assuming the node is not executing a mission;
- μ_{it} . is the minimum (relative) energy level reached by the supercapacitor of node N_i during time epoch $t, t \in \tau_i^c$, if the node is executing a mission.

Figure 4 shows an example of how the parameters Δ_{it} , Δ_{it}^I and μ_{it} are defined for a node N_i during a time epoch t . The power harvested and consumed by N_i is reported in Figure 2. In this particular example, the power consumption of the node being active and sensing is set to 12 mW. The variation of the energy level of the supercapacitor during the time epoch t is shown in Figure 4. Δ_{it}^I is computed by assuming that node N_i is not executing any mission during time epoch t . In this case, since the power harvested by N_i is greater than its power consumption, the energy stored in its supercapacitor monotonically increases (see Δ_{it}^I curve in Figure 4a). The opposite case, in which N_i is executing a mission during time epoch t , is shown in Figure 4b. In this case, the energy level of N_i 's supercapacitor decreases as long as the harvested power is lower than 12mW, reaching its minimum value μ_{it} short before 1:50PM.

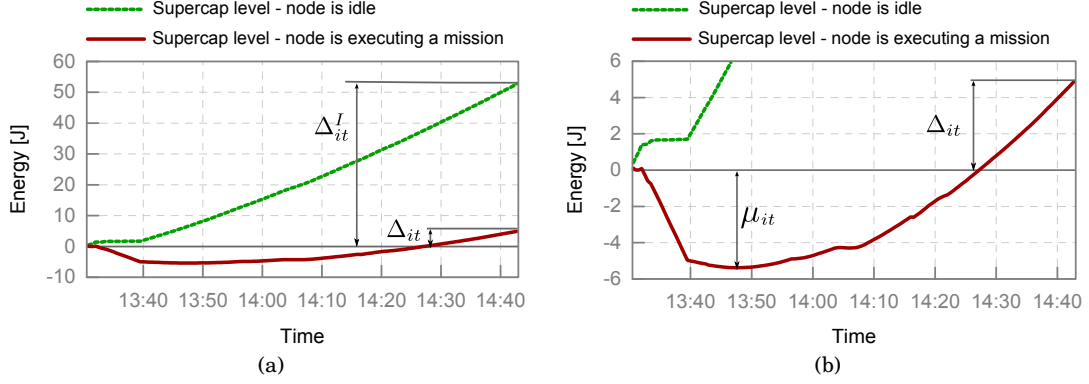


Fig. 4. (a) Variation in the energy level of the supercapacitor during time epoch t ; (b) Same figure, zoomed and annotated with the value of μ_{it} .

The parameter μ_{it} represents the minimum amount of energy that N_i must have available to execute a mission during time epoch t . As missions usually span multiple time epochs, the total energy needed to execute a mission must be computed so as to consider the μ_{it} for each time epoch. We thus define μ_{ijt} as the minimum energy level reached by the supercapacitor of node N_i if N_i starts executing M_j during time epoch t , $t \in \tau_i^c$:

$$\mu_{ijt} = \max(\mu_{it_1}, \mu_{it_2} - \Delta_{it_1}, \dots, \mu_{it_k} - (\Delta_{it_1} + \Delta_{it_2} + \dots + \Delta_{it_{k-1}})),$$

where $te_j = t + k - 1$ and, for simplification of the expression, $t_i = t + i - 1$. Thus, there are k intervals from time t to time te_j inclusive. Parameter μ_{ijt} represents the amount of energy a node must have as a reservoir to accept a given mission M_j at time t . Since μ_{it} parameters are positive (absolute) values, μ_{ijt} is defined as a maximum. In order to explain the definition of μ_{ijt} , it is useful to keep in mind that once assigned to a mission a node runs it until completion. For this reason, when computing the energy reservoir required for the node to run during the time epoch t_2 , we also take into consideration the supercapacitor energy change during the previous time epoch. For example, Δ_{it_1} represents the change in the energy level of the supercapacitor of node N_i during the time epoch t_1 . Thus, $\mu_{it_2} - \Delta_{it_1}$ is the minimum level reached by the supercapacitor of node N_i during time epoch t_2 .

For node N_i to start executing mission M_j at time t , i.e., $x_{ijt} = 1$, the energy stored in its capacitor at time t , C_{it} , plus the energy stored in its battery, B_{it} , must be at least μ_{ijt} . This is expressed by constraints 8:

$$C_{it} + B_{it} \geq \mu_{ijt}x_{ijt}, \quad \forall N_i, t \in \tau_i, j \in aM_i(t)$$

Finally, the Δ_{it} and Δ_{it}^I parameters are used to update the energy level of the supercapacitor at the end of each time epoch (constraints 5):

$$C_{it} \leq C_{it} + D_{it} + \Delta_{it}^I - (\Delta_{it}^I - \Delta_{it}) \sum_{j \in aM_i} \sum_{\substack{t' \in \tau_i: \\ ts_j \leq t' \leq t}} x_{ijt'}, \quad \forall N_i, t \in \tau_i^c$$

If the node is executing a mission, the sum $\sum_{j \in aM_i} \sum_{\substack{t' \in \tau_i: \\ ts_j \leq t' \leq t}} x_{ijt'}$ will be equal to 1 and the energy level of the supercapacitor will be updated with the energy change Δ_{it} , i.e., constraints (5) become:

$$C_{it} \leq C_{it} + D_{it} + \Delta_{it}, \quad \forall N_i, t \in \tau_i^c,$$

where the variable D_{it} is the energy drained from the battery of node N_i during time epoch t .

On the contrary, if the node is idle, the parameter Δ_{it}^I will be considered when computing the new energy level C_{it} . In this case, $\sum_{j \in aM_i} \sum_{\substack{t' \in \tau_i: \\ ts_j \leq t' \leq t}} x_{ijt'}$ is equal to zero and

constraints (5) become:

$$C_{it} \leq C_{it} + D_{it} + \Delta_{it}^I, \quad \forall N_i, t \in \tau_i^c$$

Variable D_{it} will be greater than zero if the energy change computed for time epoch t is negative (i.e., $\Delta_{it} \leq 0$ or $\Delta_{it}^I \leq 0$, depending on whether the node is executing a mission or not) and the supercapacitor can not supply the needed energy.

On the contrary, if the energy change computed for time epoch t is positive or zero, no energy will be needed from the battery, and the supercapacitor will be recharged. Objective function pressure ensures no unneeded transfer of energy between the battery and the supercapacitor, unless they come at no penalty.

Since during energy pre-computations the energy level of the supercapacitor at the beginning of each time epoch t is not known, we assume it to be 0 at the beginning of each simulation. This assumption may lead to an overestimate of the energy available to the node. In fact, if the supercapacitor is not empty, it would charge to its capacity sooner, causing excess energy to be lost once it is full. However, since our goal is to obtain an upper bound on the optimal solution, we allow for such an overestimation in order to avoid adding complexity to the model.

The formulation of the sensor-mission assignment problem that we propose can capture the idle consumption of real-life systems and some non-ideal behaviors of energy buffers, namely the fact that energy storage has a finite size, and that the charging and discharging efficiency is lower than 1. However, supercapacitor self-discharge is not modeled. In fact, including realistic models of leakage and self-discharge in the analytical formulation would significantly increase its complexity, impairing scalability. Despite that, our formulation succeeds in providing an upper bound to the optimal solution. In fact, since self-discharge negatively impacts energy availability of the nodes, the profit obtained by the optimal solution is necessarily lower than the one provided by our MIP model. We accounted, however, for supercapacitor self-discharge in our experiments and performance evaluation (Section 4), and we empirically evaluate self-discharge models in Section 5.

2.5. Sensor-mission assignment problem in energy-harvesting WSNs

As a final note, we briefly discuss hardness results for the sensor-mission assignment in energy-harvesting WSNs problem.

Bar-Noy et al. introduced the Semi-Matching with Demands (SMD), a version of the sensor-mission assignment problem in which the set of missions to be executed is known, there is no time dimension in the problem and no energy restriction limits sensor-mission assignments [Bar-Noy et al. 2008]. Specifically, an instance of SMD is a weighted bipartite graph $G = (N, M, P, E)$ where $N = \{N_1, \dots, N_n\}$ is a collection of nodes, $M = \{M_1, \dots, M_m\}$ is a collection of missions, $P = \{p_1, \dots, p_m\}$ is a collection of positive mission profits, and E is a collection of non-negative weights (utilities) for the edges $N \times M$. The demand of each mission is assumed to be 1. Missions can not be partially satisfied, i.e., $y_j \in \{0, 1\}$. The goal of SMD is to find a semi-matching $F \subseteq E$ in which $\sum_{M_j \in A} p_j$ is maximized, where $A \subseteq M$ is the set of missions satisfied by F . The corresponding decision problem is to determine, given an instance of SMD and a value b , whether a semi-matching $F \subseteq E$ such that $\sum_{M_j \in A} p_j \geq b$ exists.

Bar-Noy et al. proved that SMD is NP-hard and as hard to approximate as Maximum Independent Set. Such hardness results also apply to sensor-mission assignment

in energy-harvesting WSNs. To prove it, we give a polynomial-time reduction from a given instance of the SMD problem to an instance of the sensor-mission assignment in energy-harvesting WSNs problem.

Proposition 1 The sensor-mission assignment problem in energy-harvesting WSNs is NP-hard and at least as hard to approximate as SMD.

PROOF.

Given an instance of SMD, an instance of the sensor-mission assignment in energy-harvesting WSNs decision problem is created as follows.

The set of nodes N' of the new instance is the same as N . For each mission M_j , create a new mission M'_j with profit $p'_j = p_j$ and demand $d'_j = 1$. The duration of each mission M'_j is set to one timeslot and the start time of the new missions is set so that as each mission M'_j starts at the beginning of timeslot t_1 , i.e., missions are all overlapping. The utility e'_{ij} provided by node N'_i to mission M'_j is the same as in the SMD instance, i.e., $e'_{ij} = e_{ij}$. The mission satisfaction threshold T_{sat} is set to 1, i.e., the demand of a mission must be fully met in order for it to be satisfied. The battery energy of each node is set such as to potentially execute all the missions, i.e., $B_{max} = ec_s$, where ec_s is the energy consumed by the node in one timeslot of sensing. Finally, we assume that no energy is harvested by the nodes and that their supercapacitor is not in use, i.e., $C_{max} = 0$. This reduction requires time polynomial in the size of the SMD instance. The corresponding decision problem is to determine if there exists a sensor-mission assignment such that the total profit achieved by the network at least b .

The SMD decision problem instance has a positive answer if and only if the constructed sensor-mission assignment decision instance in energy-harvesting WSNs has a positive answer. \square

Furthermore, generalizing the analysis in [Johnson et al. 2010] to energy-harvesting wireless sensor networks, it can be proven that the online version of the sensor-mission assignment in energy-harvesting WSNs problem can not be solved with any finite competitiveness guarantee.

Proposition 2 There is no constant-competitive algorithm for the online sensor-mission assignment problem in WSN with energy harvesting, even assuming that each mission lasts only for one timeslot, that missions are not overlapping and that no energy is harvested by the nodes.

PROOF. In the following, we denote with ec_s the energy consumed by the node to perform one timeslot of sensing, while ec_i is the energy required to stay idle for one timeslot.

We first consider the case in which all nodes have only limited amount of battery energy, which allows them to execute a mission for just one timeslot and to stay idle for another timeslot, i.e., $B_{max} = ec_s + ec_i$. Since we are assuming that no energy is harvested by the nodes, the supercapacitor is empty and not in use, i.e., $C_{max} = 0$. Suppose that at time 1 a mission M_1 arrives in the network. Such mission has a profit $p_1 = \epsilon$ and necessarily requires a given sensor N_i to be assigned to it in order to be satisfied. The online deterministic assignment algorithm must assign N_i to M_1 , as there may be no further missions and refusing M_1 could lead to an infinite competitive ratio. However, at time 2 a mission M_2 may arrive in the network, having profit $p_2 = 1$ and requiring sensors N_i as well. But the remaining energy would be not enough to execute M_2 , so the total profit obtained by the online deterministic assignment algorithm would be ϵ . The optimum assignment algorithm, instead, will ignore M_1 , staying idle during timeslot 1, and then it will assign sensor N_i to M_2 when it arrives, for a profit of 1. Now suppose that the battery energy of the nodes allows them to execute a mission for two timeslots and to stay idle for another timeslot, i.e., $B_{max} = 2ec_s + ec_i$. We are

still assuming that no energy is harvested by the nodes. Then consider the following sequence of missions arrival: at time 1 missions M_1 arrives with profit $p_1 = \epsilon$, at time 2 missions M_2 arrives with profit $p_2 = \exp(\epsilon)$ and at time 3 missions M_3 arrives with profit $p_3 = \exp(\exp(\epsilon))$. M_1, M_2 and M_3 all require sensor N_i to be assigned to them in order to be satisfied. The online deterministic assignment algorithm must assign N_i to M_1 at time 1, as there may be no further missions. At time 2, N_i must be assigned to M_2 for the same reason. At time 3, there will not be enough energy left to execute M_3 , thus leading to a total profit of $\epsilon + \exp(\epsilon)$. The optimum assignment algorithm, instead, will ignore M_1 and assign sensor N_i to M_2 and M_3 when they arrive, for a profit of $\exp(\epsilon) + \exp(\exp(\epsilon))$. This construction can be extended to arbitrary B_{max} . \square

3. EN-MASSE

In this section we introduce a decentralized heuristic for sensor-mission assignment in energy-harvesting wireless sensor networks, called EN-MASSE (ENergy harvesting-aware sensor-Mission ASSignmEnt algorithm). We begin describing the communication protocol used by the nodes in the network, then we detail the behavior of our distributed scheme and the rationale for our proposed approach. The algorithm used by the mission leader to decide sensors-to-missions assignments is also discussed. Finally, we describe the energy prediction models used by EN-MASSE.

3.1. Communication protocol

Each mission arrives in the network at a specific geographic location g_j . In EN-MASSE the sensor node closest to g_j is selected as the *mission leader* and coordinates the assignment of nodes to missions. Mission leaders may be selected either by the sink in a centralized fashion or using geographic routing techniques. We selected the communication protocol described in [Johnson et al. 2010; Rowaihy et al. 2008] for exchanging information between the mission leader and the nearby nodes. We recap it here for clarity. Each time a new mission arrives in the network, the leader advertises mission information, including the mission's location, profit, demand, and duration to its neighbors. When a nearby node hears such an advertisement message, it makes a decision either to bid for the mission and become eligible for selection by the leader or to ignore the advertisement. If the node is already assigned to a mission, it ignores the advertisement message. Bidding decisions are taken autonomously by each node in the neighborhood, based on the *bidding scheme* they use (Section 3.2). Bidding nodes communicate their availability to the mission leader, which greedily selects which of the available sensor nodes to assign to the mission (Section 3.3), 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.

3.2. Bidding scheme

On the reception of a mission advertisement message sent by the leader, nodes autonomously decide to bid for participation in the mission or not. Bidding decisions are based on the following factors:

- (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 model² ;
- (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.

²See Section 3.4.1 for additional details about the solar energy prediction model we use.

A node uses the first three factors in the list to classify the incoming mission according to the impact that its execution would have on the energy reservoir of the node. More specifically, missions are classified into one of the following four classes:

Free missions. are those arriving when the node super-capacitor is full and their energy cost is expected to be fully sustained by the energy harvested during their duration. The energy prediction model estimates the amount of energy that will be harvested in the near future.

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 sustained using only the supercapacitor but this cost is not expected to be recovered through harvesting in the near future.

Battery-required missions. are those whose energy cost must be totally or partially supplied by the battery.

Such classification is used by the node to decide whether to bid for a mission or not. Nodes always accept *free* missions, as the energy they need can be directly provided by the harvesting subsystem. Ideally, executing a mission of this class should not affect the energy reservoir of the node at all. 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 provide by participating in that mission with the *expected partial profit* \bar{p} of a typical mission, which is computed based on the distribution of mission profit and demand and the expected utility contribution that a node can offer to a typical mission in its range:

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

where $E[u]$ is the expected utility contribution, $E[d]$ and $E[p]$ are the expected demand and the expected profit of a typical mission, learned by the nodes based on previous history or obtained at design time based on the application characteristics, and P is the maximum mission profit.

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

$$p^* = \frac{e_{ij}}{d_j} \times \frac{p_j}{P} \times w, \quad (17)$$

where e_{ij} is the potential utility contribution that the node N_i can provide to the given mission, d_j and p_j are, respectively, the mission demand and profit, P is the maximum mission profit and w 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 \bar{p} , to have an indication of how profitable the mission is. A node will bid for a given mission only if $p^* \geq \bar{p}$.

To decide which missions are classified as recoverable, nodes estimate the missions interarrival time and set the recoverable time interval dynamically, so that its duration is at most the expected time before the next mission arrives. In this way, choosing to

execute a recoverable mission is not likely to affect the node's ability to bid on the next mission.

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 will be provided by the supercapacitor if no energy is harvested. The same approach is taken to compute the partial profit p^* associated with the battery contribution. However in this latter case we use an additional factor when computing 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 network 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 γ that a sensor's offer is accepted. We denote with rtl 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: $e_r = rtl \times t_e \times ec_s$, where ec_s is the average mission sensing cost.

The factor w_e is then multiplied by w in the computation of p^* :

$$p^* = \frac{e_{ij}}{d_j} \times \frac{p_j}{P} \times w \times w_e,$$

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.

As a final note, we observe that, while all parameters can be easily learned or they are known a priori, the γ parameter 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 bid 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.

Algorithm 1 shows the pseudo code of the EN-MASSE bidding scheme.

3.3. Task allocation to sensors

Once the mission leader has collected bidding answers from nearby nodes, it performs the selection of nodes to assign to the mission. Nodes are ranked based on the utility they offer. The mission leader greedily assigns them to the mission, based on their ranking, until the mission is fully satisfied or there are no more sensors bidding for it. Once the minimum satisfaction threshold is reached, the leader may decide to stop short of full coverage if the next node to add would go way past the necessary utility, as there is no additional profit earned by the network if the required demand is exceeded. The mission (partially) succeeds if it reaches the success threshold T_{sat} ; if not, the leader releases all sensors. If the mission can be successfully started, an ACK message

ALGORITHM 1: EN-MASSE bidding scheme**Data:**

B_{it} . the battery energy level of node N_i at time t
 C_{it} . the supercapacitor energy level of node N_i at time t
 $M_j = (p_j, d_j, g_j, ts_j, te_j)$. the incoming mission M_j
 MT_j . the class of the mission M_j
 $EC_j(t)$. the amount of energy required to execute mission M_j from time t until its end time
 \bar{p} . the expected partial profit of a typical mission
 P . the maximum mission profit
 $e_r(t)$. the amount of energy the node expects to need to function to the target lifetime

Result: 1 if the node bids at time t for participation in the mission M_j , 0 otherwise.

if $MT_j == Free$ **then**

 | **return** 1; // Always accepts Free missions

end

// Compute the partial profit achievable by participating in M_j

$p^* = e_{ij}/d_j \times p_j/P$;

// Weight p^* based on mission's classification

switch MT_j **do**

 | **case** *Recoverable*

 | $p^* = p^* \times w_{recoverable}$; // Compute p^* , Recoverable mission

 | **case** *Capacitor-sustainable*

 | $p^* = p^* \times w_{capacitor}$; // Compute p^* , Capacitor-sustainable mission

 | **case** *Battery-required*

 | // Separately evaluate the energy contribution of the supercapacitor and the battery

 | $p_{capacitor}^* = p^* \times w_{capacitor-sustainable}$; // Compute $p_{capacitor}^*$

 | $p_{battery}^* = p^* \times w_{battery-required} \times B_{it}/e_r(t)$; // Compute $p_{battery}^*$

 | $p^* = C_{it}/EC_j(t) \times p_{capacitor}^* + (1 - C_{it}/EC_j(t)) \times p_{battery}^*$;

endsw

// Decide whether to bid for the mission, based on how profitable it is

if $p^* \geq \bar{p}$ **then**

 | **return** 1;

else

 | **return** 0;

end

is sent to the nodes chosen by the leader. When a mission ends, the leader sends out a message to announce mission termination, and all its assigned sensors are released.

The leaders of missions that have not reached their success threshold or are not fully satisfied after the first assignment process will retry to obtain more sensors when a nearby mission terminates and releases its sensors. Such information can be easily obtained by overhearing the message announcing the end of a mission.

3.4. Energy prediction models and solar energy datasets

3.4.1. Short-term solar energy predictions. EN-MASSE uses an energy prediction model to classify incoming missions, by estimating how much time will be necessary for the node to recover through harvesting the energy required for mission execution. Since such estimations are carried over a relatively small time frame (usually, half an hour to several hours), we refer to this prediction model as the short term one.

Several energy prediction models have been proposed in the literature [Kansal et al. 2007; Recas Piorno et al. 2009; Cammarano et al. 2012; Moser et al. 2007a; Noh and Kang 2011; Lu et al. 2010]. We specifically focused on three of them: the Exponential Weighted Moving Average (EWMA) prediction model [Kansal et al. 2007], the Weather-Conditioned Moving Average (WCMA) algorithm [Recas Piorno et al. 2009], and the PROfile energy (Pro-Energy) predictor [Cammarano et al. 2012]. WCMA and EWMA assume the energy generation on a typical day to be related to the energy generation at the same time on the previous days. EWMA maintains the amount of energy available during the past days as a weighted average, in which the contribution of older data decays exponentially. EWMA exploits the diurnal solar energy cycle and adapts to seasonal variations, but it exhibits significant prediction errors when weather conditions change frequently, i.e. when sunny and cloudy days are mixed. WCMA improves over EWMA by including a factor that measures the solar conditions in the present day relative to the previous days, which is especially important in frequently changing weather conditions. The main idea of Pro-Energy, instead, 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. Such energy profiles are then used to forecast the future expected energy intake: once per timeslot, Pro-Energy delivers energy predictions by looking at the stored profile that is the most similar to the current day.

We used the real-life solar data we collected (Section 4.1) to verify the accuracy of Pro-Energy, WCMA and EWMA. Based on such assessment, we selected Pro-Energy as the short-term energy prediction model used by EN-MASSE, as it consistently outperforms the other predictors. For example, on our dataset of solar traces collected in Rome, Pro-Energy exhibits a performance improvement, in terms of reduction of the mean absolute percentage error in energy predictions, of 9% – 43% with respect to EWMA and of 5% – 26% with respect to WCMA [Cammarano et al. 2012].

3.4.2. Long-term solar energy predictions. The short term energy predictor allows EN-MASSE to classify incoming missions, but it does not take into account the long-term behavior of the energy source. To account for this aspect, we propose a lightweight long-term predictor that allows to fine-tune the behavior of EN-MASSE over longer time frames. Such a predictor uses historical weather data, such as those reported in Figure 5, to obtain the average energy harvested per day during each month in which the network is expected to be operational.

The long-term predictor works as follows: at the beginning of each month, it computes a factor, LTP_m , indicating how much the average harvested energy per day is expected to change during the next month with respect to the current month. This factor allows to estimate when the energy stored in the battery of EN-MASSE nodes is more likely to be needed over the whole target lifetime, and to plan its usage accordingly. Specifically, the LTP_m factor is defined as:

$$LTP_m = \left(\frac{HE_{m+1}}{HE_m} \right)^\alpha \left(\frac{HE}{HE_m} \right)^\beta, \quad (18)$$

where:

- m . is the current month;
- HE_m . is the average energy per day harvested during the current month;
- HE_{m+1} . is the average energy per day the node expects to harvest during the next month;
- HE . is the mean of the average energy per day harvested during the whole target lifetime;
- α, β . are two weighting parameters, $a, b \in [0, 1]$.

The parameters HE_m, HE_{m+1} and HE are all computed based on historical weather data.

The first factor in Equation 18 represents how the amount of energy harvested per day is expected to change during the next month. A ratio greater than 1 indicates an increment in the average energy harvested per day, meaning that EN-MASSE would ideally use more energy from the battery during the current month than during the next one.

The second factor in Equation 18 is the ratio between the average energy harvested per day over the whole target network lifetime and the average energy harvested per day during the current month. If this ratio is bigger than 1, in the current month the average harvested energy per day is below the mean. EN-MASSE should thus be allowed to use more battery energy during the current month.

LTP_m is used as a weighting factor in Equation 17, in order to fine-tune the eagerness of EN-MASSE nodes in accepting battery-required missions. Only for this class of missions, the weight factor $w_{battery}$ is updated every month m and becomes:

$$w_{battery} = \max\{w_{battery} * LTP_m, \text{max weight}\},$$

where max weight is the maximum weight that can be assigned to a battery-required mission, which should be lower than that assigned to the other mission classes.

As the LTP_m factor computed by the long-term predictor is not used for mission classification, different long-term predictors can be seamlessly integrated into EN-MASSE, without affecting its general bidding mechanism.

4. PERFORMANCE EVALUATION

In this section we evaluate our sensor-mission assignment scheme in several different scenarios and compare its performance to that of other schemes proposed in the literature. For this purpose, we implemented a dedicated simulator in C++ that is the base of our evaluation framework for energy-harvesting WSNs. It accurately models the energy harvesting subsystem of the nodes and allows the user to import solar trace datasets in different formats.

4.1. Solar energy traces

We use real solar data to evaluate the performance of sensor-mission assignment schemes in networks with energy harvesting capabilities. Solar energy profiles were collected by using a Telos B mote [Crossbow Technology 2004] equipped with a XOB17-04x3 solar cell [IXYS Corporation 2009]. The motes were deployed in downtown Rome, for a total of 120 non-consecutive days at variable weather conditions and in different locations. A dedicated TinyOS [Levis et al. 2005] application was developed to track the amount of energy generated by the cell at 30-second intervals.

We also obtained extensive solar traces from the US Climate Reference Network [USCRN 2011], a network of climate stations deployed in the continental United States. The selected traces report the solar radiation in the period from January 2010 to December 2010, observed in four different weather stations: Santa Barbara, Boul-

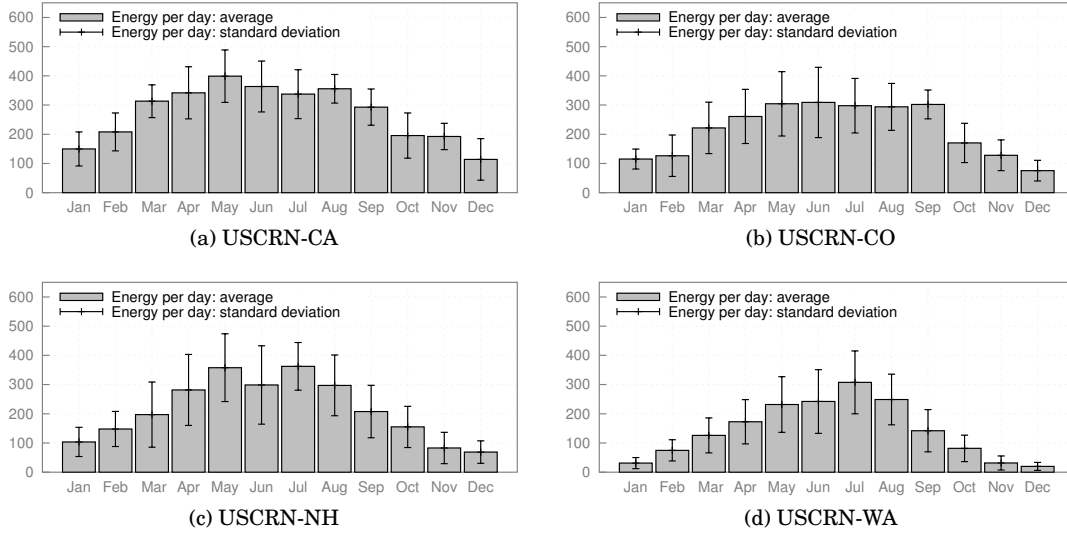


Fig. 5. Harvested energy per day: average value in Joules and standard deviation for each month and each dataset.

der, Durham and Darrington. Solar data are recorded hourly, reported as an average of the solar radiation over each hour and expressed in W/m^2 .

In order to use such data in our simulations, we scaled down the traces in each dataset by a factor equal to the ratio between the maximum peak power recorded in our solar traces and the maximum peak reported in the dataset [Vigorito et al. 2007].

Figure 5 shows how the amount of the harvested energy in the scaled datasets varies over time: for each month of the year and for each USCRN dataset, the average value and the standard deviation of the energy harvested per day is reported.

4.2. Simulation setup

In our simulations 500 nodes are randomly and uniformly scattered in a square area of $400 \times 400m$. The communication range of the nodes is set to 40m. The sensing range is 30m. The node energy model is that of ECO nodes [Park et al. 2005; Park and Chou 2006b], an ultra-compact expandable wireless sensor platform, which has been used in combination with the Ambimax energy harvesting platform [Park and Chou 2006a]. The active power consumption of the ECO node is 9mW, while its idle power consumption is 0.006mW. The sensor node is also equipped with a supercapacitor, used as an energy buffer for the energy harvested by the solar cell. The reference supercapacitor we consider is the 25F ultracapacitor of Maxwell HC Power series [Maxwell Technologies, Inc 2009a], which can nominally hold around 90J of charge. We set the charging and discharging efficiency of the supercapacitor to 95%, i.e., $\eta_c^C = \eta_d^C = 0.95$ [Barker 2002; Maxwell Technologies, Inc 2009b]. In our simulations, we model the leakage $leak_i(t)$ experienced by the supercapacitor C_i at time t by using a piecewise linear approximation of the empirical leakage pattern [Zhu et al. 2013]:

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

where $B_{R_1}, \dots, B_{R_{n+1}}$ are the residual energy values in which the slope of the leakage curve change significantly and $a_1, \dots, a_n, b_1, \dots, b_n$ are constants depending on the su-

Table II. Simulation scenarios: default settings.

Parameter	Value
Nodes in the network	500
Simulation field	400 × 400 m
Communication range	40 m
Sensing range	30 m
Average mission profit	10
Average mission demand	2
Average mission duration	1 hour
Satisfaction threshold	50% of the mission demand

percapacitor used, which represent the coefficients of the line segments used for the approximation.

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. For simplicity, we define the utility that a sensor N_i can potentially offer to a mission M_j as a function of the distance D_{ij} between the location g_j of the mission and the position of the node N_i .

Depending on the experiment, the network target lifetime has been varied between 30 and 180 days.

Table II reports the parameters used in the simulation setup.

In our implementation of EN-MASSE, we set the weighting factors w associated with mission types equal to 0.95 for battery required missions, 1.05 for capacitor sustainable missions and 1.2 for recoverable missions.

Additionally, 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 sensing range. Doing so, whenever a mission is recoverable bidding for it is not expected to compromise the node capability to serve future missions.

4.3. Benchmark assignment schemes

We compare the performance of our EN-MASSE assignment scheme with the three different energy harvesting unaware mission assignment schemes proposed in [Johnson et al. 2010], namely the Basic Scheme, the Energy Aware Scheme and the Energy-Lifetime Aware Scheme. The Energy-Lifetime Aware Scheme is shown to perform very well in traditional scenario (without harvesting) and with very limited network lifetimes (three days). In such cases, it outperforms the other assignment schemes, Basic and Energy Aware, achieving a total profit that is up to 22% higher.

Our benchmark assignment schemes are:

- **Basic Scheme:** Sensors propose to any mission within their range.
- **Energy Aware Scheme:** This scheme does not use any classification of missions or mechanism to account for the harvested energy. It also does not take the target lifetime of the network into account while taking bidding decisions. The Energy Aware Scheme evaluates how profitable a mission is by comparing the partial profit a node can achieve by participating in the incoming mission with respect to the expected partial profit of a typical mission (as defined in Section 3.2). In order to take the current energy status of a node into account, the partial profit computed for

the current mission is weighted by a factor equal to the fraction of residual battery energy available.

- **Energy-Lifetime Aware Scheme:** This scheme does not account for harvested energy, but it considers the target lifetime of the network while taking bidding decisions. It evaluates how profitable a mission is in a way similar to the Energy Aware Scheme, but it uses a different weighting factor. The partial profit computed for the current mission is weighted by a factor equal to the ratio between the time a node can be actively sensing, given its residual energy level and the target lifetime, and the expected occupancy time of the node.

Because these schemes were not initially designed for nodes with energy harvesting capabilities, we modified them so as to sum the contribution of both the capacitor and the battery when computing the residual and the maximum energy level of the node. All the bidding schemes we considered use the communication protocol described in Section 3.1 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 [Rowaihy et al. 2008].

4.4. Simulation results

We compare EN-MASSE, the Basic Scheme, the Energy Aware Scheme, and the Energy-Lifetime Aware Scheme 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.

4.4.1. Reference scenario. Our first set of experiments refer to a scenario where node sensing cost is 4.5 mW, e.g., the power consumption of an ultra low-power vision sensor such as [Gasparini et al. 2011]. Each node is equipped with two solar cells and with a 25F Maxwell supercapacitor. 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 20 missions per hour. The solar dataset we use is that of real-life solar traces collected by using Telos B mote and XOB17-04x3 solar cells (Section 4.1). All other parameters are as detailed above.

We evaluate the performance of the different protocols with respect to the following metrics:

- (1) profit achieved over time by the network (Figure 6a);
- (2) total profit achieved at the network target lifetime (Figure 6b);
- (3) profit achieved per day: mean value and standard deviation (Figure 6c);
- (4) fraction of nodes that have energy left in their battery over time (Figure 6d).

The first three 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. In particular, the third metric provides an indication of the stability and robustness of the assignment scheme. The fourth metric deserves some discussions. 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, and it may be 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

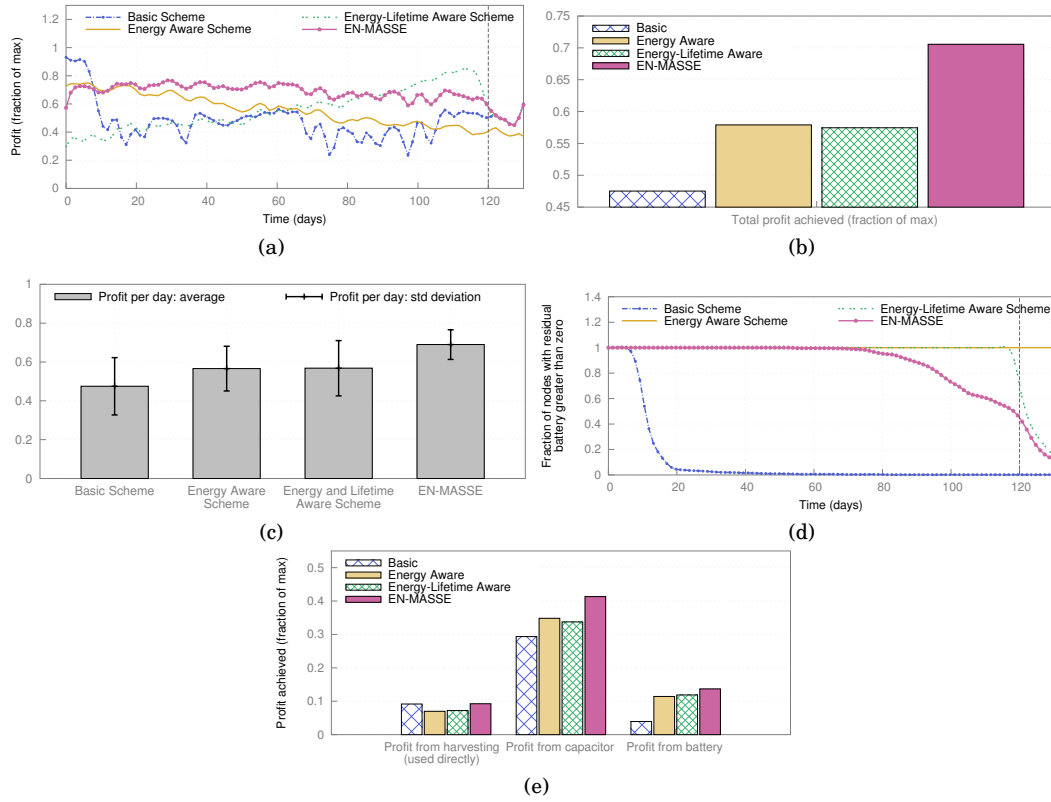


Fig. 6. Simulation results for a sample scenario [Rome scenario, real-life solar traces dataset]: (a) fraction of achieved profit per time slot, (b) total profit achieved at target lifetime (as a fraction of maximum), (c) profit achieved per day: average value and standard deviation, (d) fraction of nodes with a residual battery greater than zero and (e) total profit achieved per energy source (as a fraction of maximum).

subsystem may compromise the capability of the network to execute critical missions, degrading the network profit.

We observe that EN-MASSE leads to the greatest total profit (Fig. 6b), and that is able to provide the highest profit (70 – 80% of the maximum profit, defined as the profit that could be achieved if all missions were fully served) till the target network lifetime (Fig. 6a). Moreover, as shown in Fig. 6c, EN-MASSE also provides the most stable profit over time among the considered bidding schemes, as the average profit it obtains per day is the least variable. The fact nodes wisely exploit their resources to provide the needed support to arriving missions is confirmed by Fig. 6d, which shows that no node has depleted its battery in the first 70 days. Then the fraction of nodes with no energy in their battery starts increasing, reaching 45% 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 deplete their battery quickly, 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 nodes' batteries at the target network lifetime is wasted. Fig. 6e 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 (Fig. 6c), 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. 6b, 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, by becoming more conservative as the fraction of the residual battery energy decreases. This trend is clearly visible in Figure 6a, which shows as the profit achieved by the Energy Aware Scheme decreases over time. Thanks to this mechanism, as can be seen in Fig. 6d, no node has completely depleted its battery after 120 days of simulation. However, to achieve this the Energy Aware 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 as the target network lifetime approaches, and increasing its profit. For this reason, this scheme is not able to provide a stable profit over time: as shown in Fig. 6c, the variability of the profit it achieves per day is the highest among all the bidding schemes. Overall, the total profit achieved during the target network lifetime is quite low: 57% the maximum and 20% lower than what achieved by EN-MASSE.

4.4.2. Impact of the long-term energy predictor. In this section we discuss the impact of the long-term energy predictor on the performance of EN-MASSE. We focus on the solar dataset collected in Darrington, WA, from September to December 2010. We use historical data of weather conditions in the past years, which are available for this scenario, to analyze the impact of the long-term energy predictor (LTP) described in Section 3.4.2 on the performance of EN-MASSE. In the setup we consider, all the simulation parameters (network topology, mission arrival rate, target lifetime, etc) are the same as described in Section 4.4.1.

Fig. 7 shows the comparison of the performance obtained by EN-MASSE with and without the use of the long-term predictor. As can be seen in Fig. 7a, EN-MASSE without LTP, despite still outperforming the other assignment schemes, is not quite able to provide a stable profit over time, as its performance degrades as the target network lifetime approaches. This is due to the fact that, without using the long-term predictor, EN-MASSE tends to consume battery energy at approximately the same rate over time (Fig. 7c). However, as shown in Figure 5d, in the four months during which the network is required to be operational the average energy harvested per

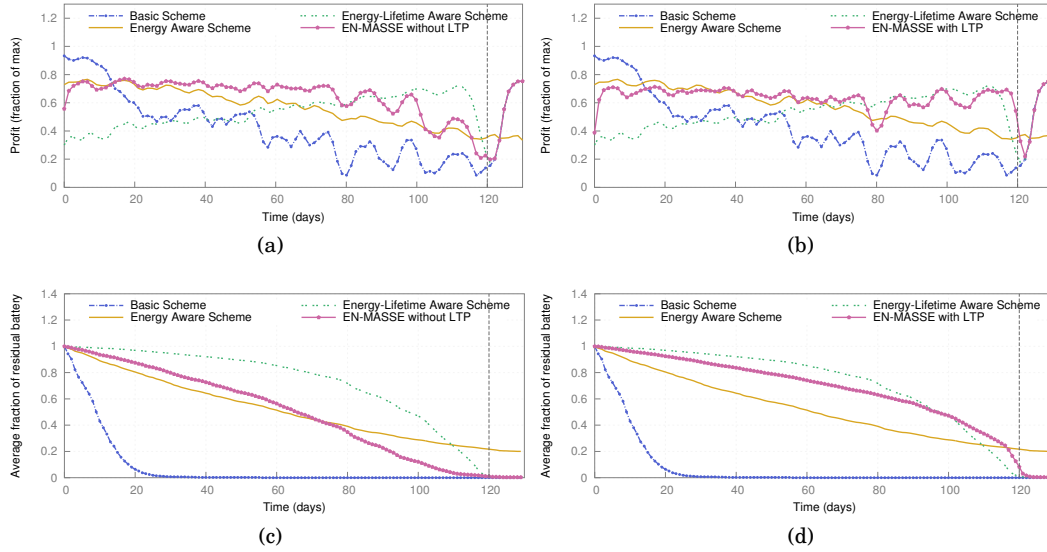


Fig. 7. Impact of the long-term energy predictor on the performance of EN-MASSE [Darrington solar traces dataset, September-December 2010]: fraction of achieved profit per time slot without (a) and with (b) LTP; average residual battery energy over time without (c) and with (d) LTP.

day significantly and steadily decreases, passing from 142 J in September to 20 J in December.

In EN-MASSE with LTP the long-term predictor is used to estimate how the harvesting trends are going to change from month to month, in order to decide when the energy stored in the battery is more likely to be needed. The effectiveness of this approach is confirmed by Figure 7b, which shows how EN-MASSE with LTP is able to achieve a much stabler profit over time, by consuming energy from the battery at different rates in different months (Fig. 7d).

4.4.3. Impact of parameter variations on performance. In this section we investigate how varying key parameters affect the performance of EN-MASSE and of the other considered mission assignment schemes. The parameters that have been varied is indicated in the title of the following subsections. Unless otherwise specified, scenario and protocol parameters are the same detailed in Section 4.2. Fig. 8 shows the ratio between the total profit achieved at target lifetime by EN-MASSE and the profit achieved by the other schemes. Each data point is obtained by averaging the results over 10 runs; for each run, the simulation is executed until the target lifetime of the network is reached. Due to space constraints, for each scenario we do not show the behavior of each scheme over time (as in Section 4.4.1), but only the performance improvement obtained by EN-MASSE with respect to the other assignment schemes.

Impact of target lifetime. As shown in Fig. 8a, 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

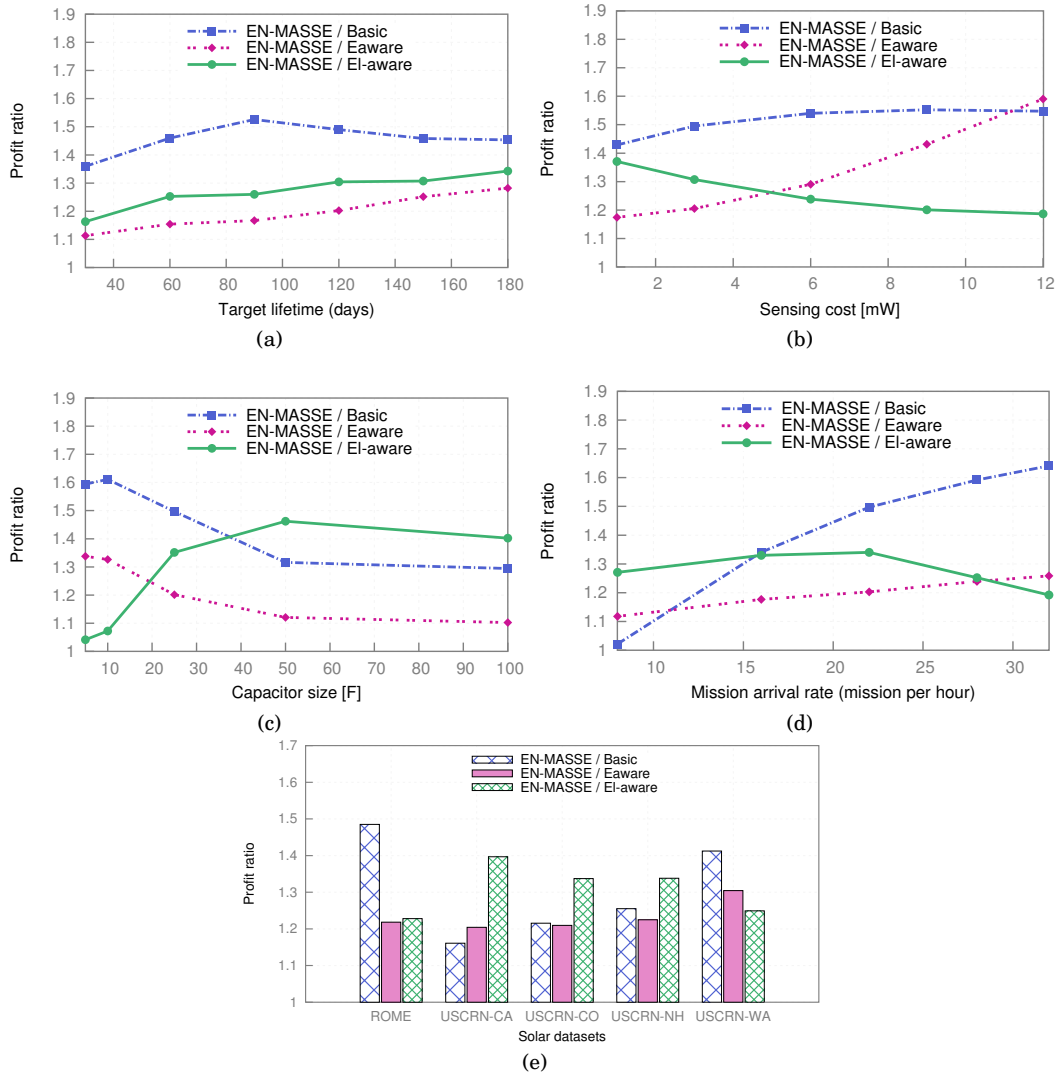


Fig. 8. Simulation results: performance improvement for varying (a) target lifetime, (b) sensing cost, (c) capacitor size, (d) mission arrival rate and (e) harvesting scenarios.

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 sus-

tained by energy harvesting, reducing the ratio between the profits achieved by the two schemes. However, the gap between the Basic Scheme and EN-MASSE remains significant. When considering target lifetimes of one year or longer, the Basic Scheme, despite outperforming the Energy-Aware and Energy-Lifetime-Aware schemes, achieves profits that are 10 – 15% lower than EN-MASSE.

Impact of sensing cost. Fig. 8b 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 self-discharge 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 that, EN-MASSE still outperforms the Energy-Lifetime Aware Scheme, achieving a profit ratios between 1.35 (lower sensing cost) and 1.2 (higher sensing cost).

Impact of capacitor size. As shown in Fig. 8c, 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 bidding 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 a 50F capacitor, because of its higher self-discharging rate.

Impact of mission arrival rate. Fig. 8d 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.

Impact of the harvesting scenario. Fig. 8e shows how the performance of EN-MASSE varies with respect to the other schemes when considering different harvesting scenarios, i.e., different solar datasets. For this set of experiments, we used the solar energy traces we collected in Rome and the ones from the US Climate Reference Network (Section 4.1). The first four months of solar data from each dataset were used.

As can be seen from the figure, our scheme, EN-MASSE, consistently outperforms the other assignment schemes by achieving a profit that is 16 – 49% higher.

The Basic Scheme is the scheme that depends the most on the harvested energy. In fact, by comparing the results shown in Fig. 8e with the statistics of the solar datasets reported in Fig. 5, it can be clearly seen that its performance increases for scenarios in which the average harvested energy per day is higher.

The performance of the Energy-Lifetime Aware Scheme also varies for different harvesting scenarios. This scheme does not exploit information about the harvested energy. For this reason, the higher the average harvested energy per day is, the more pronounced is the gap between the performance achieved by the Energy-Lifetime Aware Scheme and EN-MASSE.

The performance of the Energy Aware Scheme remains quite constant over all the different harvesting scenarios, but it exhibits the worst performance in the USCRN-WA scenario. This is due to the fact that, in this scenario, the energy harvested in the first couple of months is quite low. When running the Energy Aware scheme, nodes in the network have, on average, depleted more than 50% of their battery within the first 40 days. Since this scheme becomes more conservative as the fraction of the residual battery energy decreases, the Energy Aware Scheme starts refusing many missions and it is not able to take advantage of the fact that the average energy harvested per day increases in the last two months.

4.5. Comparison with the MIP solution

We compare our solution to the optimal mixed integer programming (MIP) solution, obtained by solving the formulation of the problem described in Section 2.4.1. Solving this model provides an upper bound on the optimal solution for the sensor-mission assignment in wireless sensor networks with energy harvesting.

We consider a network of 200 nodes, deployed in a field of 400×400 m; the arrival rate of missions in the network is 4 missions per hour. The battery size has been scaled accordingly, since the problem is less interesting if there is enough energy to execute almost all missions; the harvested energy and the power consumption of the node, both in idle mode and when executing a mission, are the same as in the previous sections. To ease the computation, the satisfaction threshold, T_{sat} , has been set to zero for these instances.

In order to remain consistent with the problem formulation reported in Section 2.4, we consider in this section a slightly modified problem. In fact, in our MIP model problem constraints 8 impose that a node is only allowed to accept a mission if it has enough energy to run until completion. In the simulation results previously reported, instead, such constraints were not enforced. We have performed extensive simulations also for this more constrained scenario, which confirm the same trend discussed in Section 4.4. Fig. 9 shows the performance of EN-MASSE with respect to the upper bound provided by the optimal MIP solution for different target lifetimes, sensing costs, capacitor sizes and mission arrival rates. The y -axis shows the total profit obtained at target lifetime as a fraction of the maximum profit. The gap between our solution and the upper bound provided by our MIP model varies between 0.54% and 11.29% of the maximum profit, depending on the considered scenario, with an average gap of 5.55% of the maximum profit.




Fig. 9. Performance of EN-MASSE with respect to the upper bound provided by our MIP models with ideal supercapacitors (MIP ideal sc) and with supercapacitor non-idealities (MIP non-ideal sc) for varying (a) target lifetime, (b) sensing cost, (c) capacitor size and (d) mission arrival rate.

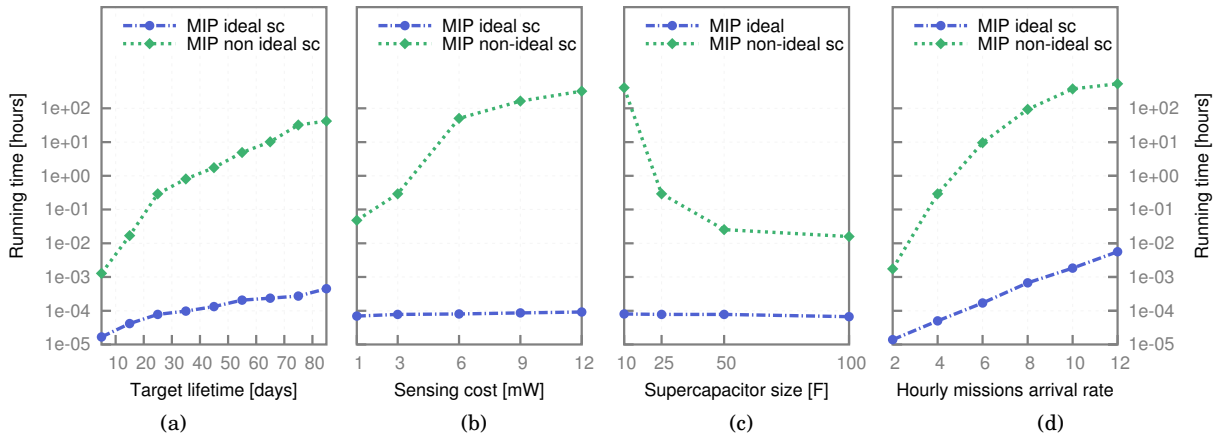


Fig. 10. Running times in hours of our MIP model with ideal supercapacitors and of our MIP model with supercapacitor non-idealities for varying (a) target lifetime, (b) sensing cost, (c) capacitor size and (d) mission arrival rate.

We also compare the upper bound provided by our MIP model, taking into account the non-idealities of supercapacitors (Section 2.4.1), with the upper bound provided by a similar formulation, which instead models *ideal* supercapacitors. An ideal supercapacitor stores any amount of energy without inefficiencies in charging and discharging, so that any excess harvested energy is stored and later used without losses. The MIP formulation of such ideal case is the same as detailed in Section 2.4.1, with C_{max} being unbounded and $\eta_c^C = \eta_d^C = 1$.

As shown in Figure 9, the gap between our model, which takes into account supercapacitors non-idealities, and the ideal model varies between 0.05% and 17.13% of the maximum profit, depending on the considered scenario.

We solved the MIP using the CPLEX commercial MIP solver version 12.3 on a Linux machine with 8 cores and 16 GB of RAM. Figure 10 shows the running times in hours of

our MIP models for varying target lifetimes, sensing costs, capacitor sizes and mission arrival rates³. The running times of both our MIP model with supercapacitor non-idealities and of our MIP model with ideal supercapacitors are reported.

As shown in the figure, the running time of the model with ideal supercapacitor is significantly lower than that of the model with supercapacitor non-idealities. This is due to the fact that considering a supercapacitor with unbounded maximum size strongly eases the computation. In fact, since no energy may be lost due to the capacitor being full, it is no longer necessary to constantly keep track of the current energy level of the supercapacitor. To enforce the energy constraints in such a case, it is enough to ensure that, at any time t , the energy spent by each node is always equal or less than the sum of the energy harvested so far plus the energy initially stored in the battery.

The comparison between the ideal model and the non-ideal model shown in Figure 9, however, highlights the importance of considering buffer non-idealities when modeling energy-harvesting systems. In fact, using a detailed formulation of realistic energy storages greatly reduces overestimation of the available energy, providing tighter bounds on the optimal solution with respect to the ideal case. Overall, the MIP model with ideal supercapacitors may be used to quickly obtain loose bounds and rough benchmark results on the achievable performance of a sensor-mission assignment algorithm.

5. EXPERIMENTAL TESTBED: IMPLEMENTATION AND VALIDATION

In order to validate our approach in a real-life application scenario, we interfaced six Telos B motes [Crossbow Technology 2004] with directional video sensors (4D Systems μ CAM [4D Systems 2011] and Link Sprite Y201 [LinkSprite 2009] color cameras). Nodes power their camera on and turn them off by using a dedicated MOS switch controlled through the GPIO port located on the Telos B primary expansion connector. The power consumption of the cameras we selected is of around 200 mW in ON mode.

We powered each node by using a harvesting system composed of a 0.5W solar panel, two 50F Panasonic Gold supercapacitors [Panasonic 2008] and two AA primary batteries (Fig. 11). The two supercapacitors were wired in series to get a higher operating voltage, thus providing a total capacitance of 25F. The harvesting subsystem also implements a maximum power point controller (MPPC), which dynamically maximizes the harvesting efficiency, in order to get as much power as possible from the solar panel under any lighting condition. The voltage of the solar panel, the supercapacitor and of the batteries are periodically sampled and recorded by the nodes by using dedicated test points connected to the ADC input ports of the Telos B mote.

To validate our bidding scheme, we implemented EN-MASSE in TinyOS, together with a state-of-the-art energy prediction model, Pro-Energy, which is able to deliver short-term solar predictions. In more detail, we implemented the communication protocol outlined in Section 3.1, the EN-MASSE bidding scheme described in Section 3.2 and the missions allocation algorithm defined in Section 3.3, in TinyOS. To reduce the energy consumption of the communication activity, the CC2420 low power listening radio stack [Moss et al. 2007] was used.

In order to realize a fully functional system, we also implemented in TinyOS both a basic driver for the Link Sprite and the μ CAM cameras, and a simple battery state-of-charge estimator, based on a direct voltage look up table.

³MIP is NP-hard in general, but commercial solvers have become so good that these solvers can quickly and reliably solve problems with millions of variables, depending on the structure. MIP has therefore become a workhorse technology for operations-research professionals solving difficult problems for industry [Bixby and Rothberg 2007].




Fig. 11. Telos B mote interfaced with a μ CAM camera and powered by an hybrid harvesting subsystem composed of a solar panel, two supercapacitors and two AA primary batteries.




Fig. 12. Solar data collected over three days of testbed run by three different nodes.




Fig. 13. Snapshot of 9 hours of solar cell, supercapacitor and battery voltage measurements.

We deployed our testbed outdoors, on the roof terrace of University of Rome "La Sapienza" CS Dept. building, which is located in downtown Rome. The test scenario we considered is that of a network deployed for intrusion detection. Missions were generated at random locations. The utility that a node was able to offer to a given mission was defined as a function of the distance between the camera of the node and the mission location. Each node assigned to a mission was required to turn its camera on and to periodically collect pictures until the end of the mission.

The data obtained while running the testbed have been used to validate the performance trends found in simulations. Moreover, for the whole test duration we also collected data about the harvested energy, the supercapacitor charge and discharge patterns, and the battery utilization, in order to validate our energy model.

5.1. Experimental validation of EN-MASSE and of the energy models

We deployed a testbed of six Telos B motes on Nov 8, 2012 at 02:40 PM. Five of such motes were sensing nodes equipped with video cameras, while the sixth node was the sink. Missions arrived in the network at random locations with an arrival rate of five missions per hour. The average duration of a mission was set to ten minutes and the target lifetime of the network was of two weeks.

After two weeks, we stopped the test and retrieved the logged data, which were then used to validate EN-MASSE simulations against the real deployment. To accurately reproduce the actual weather conditions experienced by the deployed nodes, we used in our simulations the solar data collected by the testbed. Figure 12 shows an example of solar data recorded by three different nodes over three days.

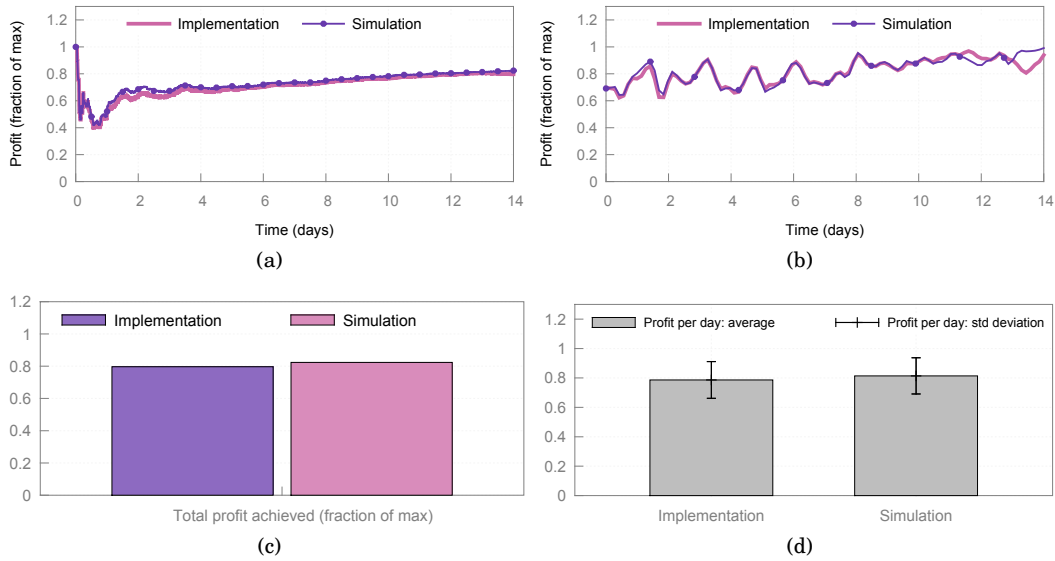


Fig. 14. Comparison between experimental performance and simulation results: (a) *cumulative* profit achieved over time; (b) fraction of achieved profit per time slot; (c) total profit achieved at target lifetime (as a fraction of maximum) and (d) profit achieved per day: average value and standard deviation.




Fig. 15. Energy model validation, comparison between measured data and simulation traces: (a) supercapacitor voltage and (b) battery voltage.

Fig. 13 shows a snapshot of 9 hours of solar cell, supercapacitor and battery voltage measurements, which highlights how different energy sources were used to power the nodes. Gray-shadowed areas in the figure represent periods of time during which the node was executing a mission.

The comparison between EN-MASSE simulations results and real-life testbed traces is shown in Figure 14.

As can be seen from the figure, simulation results are remarkably close to experimental data. In fact, the difference between the total profit computed by simulation and the profit actually achieved by the testbed is less than 3% of the maximum profit. Such gap is due to the differences between the actual energy availability of the nodes and the estimates produced by the energy models we used in our simulations.

To further investigate this aspect, Figure 15 shows the comparison of our supercapacitor and battery models with the actual voltage measurements recorded during the fourteen days of our experiment. Empirical supercapacitor data, shown in Fig. 15a, confirms the accuracy of our model, which is able to capture the changes in the supercapacitor voltages that occur during supercapacitor charging and discharging, and due




Fig. 16. Supercapacitors self-discharge: comparison between empirical discharge pattern and different leakage models.

to self-discharge redistribution. A comparison between real-life battery voltage measurements and the data obtained through simulation is shown in Fig. 15b. We found our battery model to be accurate enough to support mission bidding decisions, as it succeeds in reflecting the overall discharge patterns of the battery. However, our simple model does not capture the fine-grained behavior of battery discharge; dedicated models [Rao et al. 2003] should be employed if more accurate simulations are needed.

Finally, figure 16 shows the comparison between the empirical discharge pattern of two 50F Gold Cap supercapacitors wired in series, and their self-discharge simulated according to three different leakage models proposed in the literature.

To obtain the empirical data, we charged the supercapacitors to 2.3V for one hour, then disconnected them from the circuit. We collected self-discharge data for 17 days, by measuring the supercapacitors voltage drop once per minute.

Using such experimental data, we evaluated the accuracy of three different methods, in which the leakage experienced by the charged supercapacitors is modeled:

- (1) as a constant current [Kansal et al. 2007];
- (2) as an exponential function of the supercapacitor voltage [Renner et al. 2009];
- (3) by using a piecewise linear approximation of the empirical leakage pattern [Zhu et al. 2013].

The parameters of all the models were set as to minimize the error between estimation and actual data. As can be seen in the figure, the constant current model fails to capture the empirical discharge pattern of the supercapacitor, which is clearly not linear. The accuracy of the exponential model, instead, is quite good: its root mean squared error is around 24 mV. Finally, using the piecewise linear approximation model yields the best results, as its root mean squared error is less than 3.5 mV, i.e., one seventh of that of the exponential model. This motivates the use of piecewise linear approximation in our simulations.

6. RELATED WORK

6.1. Sensor-mission assignment problem

The general problem of sensor-mission assignment, being a key issue in the intelligence, surveillance, and reconnaissance (ISR) domain [Gomez et al. 2008; Preece et al. 2008; Le et al. 2009], has received considerable attention lately. Bar-Noy et al. introduced the Semi-Matching with Demands (SMD) in [Bar-Noy et al. 2008]. Their 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 [Rowaihy et al. 2008; Rowaihy et al. 2010], 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 [Johnson et al. 2010]. 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. Pizzocaro et al. introduced the Sensor Utility Maximization (SUM) problem in [Pizzocaro et al. 2008], in which missions are associated with uncertain demands for sensing resource capabilities. In [Le et al. 2009], Le et al. consider missions that 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. Their work was extended in [Le et al. 2010], by adding an adaptive negotiation mechanism in the allocation process. In [Rowaihy et al. 2009], Rowaihy et al. describes two utility models where the contribution provided by multiple sensors to a mission may be combined not only additively, as in previous works, but also in more involved ways. They also introduce the concept of fuzzy location, which may be useful to preserve location privacy. Their work was later extended in [Rowaihy 2012], by turning the focus on maximizing the profit achieved by the network while preserving sensors location privacy. Erol-Kantarci et al. considers in [Erol-Kantarci and Mouftah 2012] a network of wireless rechargeable sensors powered by Radio Frequency (RF)-based energy transfer. Their goal is to optimize the placement of RF-based chargers such as that the nodes with replenished batteries participate in profit maximizing missions.

6.2. Power management and task scheduling 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*). Some examples of such prototypes include [Magno et al. 2012; Carli et al. 2011; Brunelli 2008; Corke et al. 2007; Park and Chou 2006a; Simjee and Chou 2006; Raghunathan et al. 2005; Jiang et al. 2005]. Environmentally-powered systems have the potential for unlimited lifetime, but the great variability of environmental energy sources and their unpredictable nature poses new challenges in terms of energy management. New harvesting-aware power management and task scheduling techniques are thus required for network performance, lifetime enhancement and full exploitation of energy recharging opportunities.

Power management policies that allow rechargeable nodes to achieve near perpetual lifetime have been proposed in [Kansal et al. 2007; Vigorito et al. 2007]. The main idea of such works is to dynamically adapt the duty cycle of the nodes according to the environmental power conditions, in order to achieve the so called *energy-neutral* (ENO) mode, in which nodes consume only as much energy as harvested. Although effective in many applications, such duty-cycle-based approach, however, is not suffi-

cient for complex system in which more sophisticated task scheduling techniques are required.

The problem of how to schedule local tasks of nodes while dealing with uncertainty in energy availability is an active area of research, which has been addressed by many works [Zhu et al. 2012; EL Ghor et al. 2011; Audet et al. 2011; Liu et al. 2012; Recas Pirono et al. 2010; Ravinagarajan et al. 2010; Steck and Rosing 2009; Moser et al. 2007b]. For example, EL Ghor et al. describe a variant of the Earliest Deadline First (EDF) algorithm [EL Ghor et al. 2011] that runs in ENO mode. Steck and Rosing presented two adaptive algorithms [Steck and Rosing 2009] to balance task utility and execution time subject to an energy constraint aimed at guarantying energy neutrality. The interested reader is referred to [Basagni et al. 2013] for a more complete survey on task scheduling in WSNs with energy harvesting.

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. Thus, requesting the nodes to achieve energy neutrality may conflict with our goal, because some profitable missions may be rejected when the current harvesting rate is low. For this reason, we may want to violate energy neutrality when necessary to better serve important missions. Second, the problem we address is harvesting-aware *collective* task allocation at network level, i.e., how to assign nodes in the network to competitive tasks that may require the simultaneous contribution of multiple sensors. Such problem is quite different from the task scheduling problem tackled by other works, which instead focus on *individual* tasks scheduling at node level.

7. CONCLUSIONS

In this paper we have presented an analytical model and a distributed solution, EN-MASSE, for sensor-mission assignment in mission-centric WSNs with energy harvesting. We proposed a MIP formulation that captures advanced real-life details of the behavior of a typical energy harvesting subsystem, including supercapacitor and battery non-idealities, possibility to use the harvested energy directly when the current power consumption is higher than the harvested power, hybrid storage systems, and variations in the power consumption of the node. We demonstrated that considering buffer non-idealities provides a tighter bound on the optimal solution, by reducing overestimations of the available energy. Moreover, by addressing scalability issues of previous models, we shown that our formulation is able to provide upper bounds on the optimal solution for real-life sized instances, consisting in hundreds of nodes and with target lifetimes of several months. Despite of the theoretical difficulty of the general sensor-mission assignment problem, our distributed scheme, EN-MASSE, is shown to perform very closely to the optimum provided by the analytical formulation. Furthermore, EN-MASSE significantly outperforms other harvesting-unaware assignment schemes. 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. Finally, we experimentally validated our proposed scheme and the accuracy of the energy models we used in simulations in a testbed of Telos B motes powered by energy harvesting.

REFERENCES

- 4D Systems. 2011. μ Cam Serial JPEG Camera Module Data Sheet. (2011). Retrieved Jul 27, 2013 from <http://www.4dsystems.com.au>.
- David Audet, Leandro Collares de Oliveira, Neil MacMillan, Dimitri Marinakis, and Kui Wu. 2011. Scheduling recurring tasks in energy harvesting sensors. In *Proceedings of IEEE IN-*

- FOCOM 2011, Workshop on Green Communication and Networking*. Shanghai, China, 277–282. DOI: <http://dx.doi.org/10.1109/INFCOMW.2011.5928823>
- Amotz Bar-Noy, Theodore Brown, Matthew P. Johnson, Thomas La Porta, Ou Liu, and Hosam Rowaihy. 2008. Assigning sensors to missions with demands. In *Proceedings of the 3rd international conference on Algorithmic aspects of wireless sensor networks (ALGOSENSORS 2007)*. Springer-Verlag, Berlin, Heidelberg, 114–125. DOI: http://dx.doi.org/doi:10.1007/978-3-540-77871-4_11
- Philip P. Barker. 2002. Ultracapacitors for use in power quality and distributed resource applications. In *IEEE Power Engineering Society Summer Meeting, 2002*, Vol. 1. Chicago, IL, 316–320. DOI: <http://dx.doi.org/10.1109/PSS.2002.1043241>
- Novella Bartolini, Tiziana Calamoneri, Thomas La Porta, Chiara Petrioli, and Simone Silvestri. 2012. Sensor activation and radius adaptation (SARA) in heterogeneous sensor networks. *ACM Trans. Sen. Netw.* 8, 3, Article 24 (Aug. 2012), 34 pages. DOI: <http://dx.doi.org/10.1145/2240092.2240098>
- Stefano Basagni, M. Yousof Naderi, Chiara Petrioli, and Dora Spenza. 2013. Wireless Sensor Networks with Energy Harvesting. In *Mobile Ad Hoc Networking: The Cutting Edge Directions*. John Wiley and Sons, Inc., Hoboken, NJ, Chapter 20, 701–736. DOI: <http://dx.doi.org/doi:10.1002/9781118511305.ch20>
- Fang Bian, David Kempe, and Ramesh Govindan. 2006. Utility based sensor selection. In *Proceedings of the 5th international conference on Information processing in sensor networks (IPSN 2006)*. Nashville, TN, 11–18. DOI: <http://dx.doi.org/doi:10.1109/IPSN.2006.244032>
- Giuseppe Bianchi, Angelo T. Caposelle, Chiara Petrioli, and Dora Spenza. 2013. AGREE: exploiting energy harvesting to support data-centric access control in WSNs. *Ad Hoc Networks* 11, 8 (2013), 2625–2636. DOI: <http://dx.doi.org/10.1016/j.adhoc.2013.03.013>
- Robert Bixby and Edward Rothberg. 2007. Progress in computational mixed integer programming – A look back from the other side of the tipping point. *Annals of Operations Research* 149 (2007), 37–41. DOI: <http://dx.doi.org/10.1007/s10479-006-0091-y>
- Davide Brunelli. 2008. Miniaturized Solar Scavengers for Ultra-low Power Wireless Sensor Nodes. In *Workshop on Energy in Wireless Sensor Networks, WEWSN 2008*. Santorini Island, Greece.
- Alessandro Cammarano, Chiara Petrioli, and Dora Spenza. 2012. Pro-Energy: a novel energy prediction model for solar and wind energy harvesting WSNs. In *Proceedings of the 9th IEEE International Conference on Mobile Ad hoc and Sensor Systems (IEEE MASS 2012)*. Las Vegas, Nevada, 75–83. DOI: <http://dx.doi.org/10.1109/MASS.2012.6502504>
- Davide Carli, Davide Brunelli, Luca Benini, and Massimiliano Ruggeri. 2011. An effective multi-source energy harvester for low power applications. In *Proceedings of the Design, Automation and Test in Europe Conference (DATE 2011)*. Grenoble, France, 1–6. DOI: <http://dx.doi.org/10.1109/DATE.2011.5763142>
- Zichong Chen, Guillermo Barrenetxea, and Martin Vetterli. 2012. Share Risk and Energy: Sampling and Communication Strategies for Multi-Camera Wireless Monitoring Networks. In *Proceedings of the 31st Annual IEEE International Conference on Computer Communications (INFOCOM 2012)*. Orlando, FL, 1862–1870. DOI: <http://dx.doi.org/doi:10.1109/INFCOM.2012.6195561>
- Peter Corke, Philip Valencia, Pavan Sikka, Tim Wark, and Les Overs. 2007. Long-duration solar-powered wireless sensor networks. In *Proceedings of the 4th workshop on Embedded networked sensors (EmNets 2007)*. New York, NY, USA, 33–37. DOI: <http://dx.doi.org/10.1145/1278972.1278980>
- Crossbow Technology. 2004. TelosB mote platform datasheet. (2004). Document Part Number: 6020-0094-01 Rev B.
- Hussein EL Ghor, Maryline Chetto, and Rafic Hage Chehade. 2011. A real-time scheduling framework for embedded systems with environmental energy harvesting. *Computers and Electrical Engineering Journal* 37, 4 (July 2011), 498–510. DOI: <http://dx.doi.org/10.1016/j.compeleceng.2011.05.003>
- Melike Erol-Kantarci and Hussein T. Mouftah. 2012. Mission-aware placement of RF-based power transmitters in wireless sensor networks. In *IEEE Symposium on Computers and Communications (ISCC 2012)*. Cappadocia, Turkey, 12–17. DOI: <http://dx.doi.org/10.1109/ISCC.2012.6249261>
- Leonardo Gasparini, Roberto Manduchi, and Massimo Gottardi. 2011. An Ultralow-Power Wireless Camera Node: Development and Performance Analysis. *IEEE Transactions on Instrumentation and Measurement* 60, 12 (2011), 3824–3832. DOI: <http://dx.doi.org/10.1109/TIM.2011.2147630>
- Mario Gomez, Alun Preece, Matthew P. Johnson, Geeth Mel, Wamberto Vasconcelos, Christopher Gibson, Amotz Bar-Noy, Konrad Borowiecki, Thomas La Porta, Diego Pizzocaro, Hosam Rowaihy, Gavin Pearson, and Tien Pham. 2008. An Ontology-Centric Approach to Sensor-Mission Assignment. In *Proceedings of the 16th international conference on Knowledge Engineering: Practice and Patterns (EKAW 2008)*. Springer-Verlag, Berlin, Heidelberg, 347–363. DOI: http://dx.doi.org/10.1007/978-3-540-87696-0_30
- IXYS Corporation. 2009. XOB17 IXOLAR High Efficiency Solar Bits Technical Information. (Jun 2009).
- Xiaofan Jiang, Joseph Polastre, and David Culler. 2005. Perpetual environmentally powered sensor networks. In *Proceedings of the 4th international symposium on Information pro-*

- cessing in sensor networks (IPSN '05). IEEE Press, Piscataway, NJ, USA, Article 65. DOI: <http://dx.doi.org/10.1109/IPSN.2005.1440974>
- Matthew P. Johnson, Hosam Rowaihy, Diego Pizzocaro, Amotz Bar-Noy, Stuart Chalmers, Thomas La Porta, and Alune Preece. 2010. Sensor-Mission Assignment in Constrained Environments. *IEEE Trans. Parallel Distrib. Syst.* 21 (November 2010), 1692–1705. DOI: <http://dx.doi.org/doi:10.1109/TPDS.2010.36>
- Aman Kansal, Jason Hsu, Sadaf Zahedi, and Mani B. Srivastava. 2007. Power management in energy harvesting sensor networks. *ACM Trans. Embed. Comput. Syst.* 6, 4 (2007), 32. DOI: <http://dx.doi.org/10.1145/1274858.1274870>
- Thomas La Porta, Chiara Petrioli, and Dora Spenza. 2011. Sensor-mission assignment in wireless sensor networks with energy harvesting. In *Proceedings of the 8th IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (IEEE SECON 2011)*. Salt Lake City, UT, 413–421. DOI: <http://dx.doi.org/10.1109/SAHCN.2011.5984925>
- Thao Le, Timothy J. Norman, and Wamberto Vasconcelos. 2009. Agent-based Sensor-Mission Assignment for Tasks Sharing Assets. In *Third International Workshop on Agent Technology for Sensor Networks*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC.
- Thao P. Le, Timothy J. Norman, and Wamberto Vasconcelos. 2010. Adaptive negotiation in managing wireless sensor networks. In *Proceedings of the 13th international conference on Principles and Practice of Multi-Agent Systems (PRIMA 2010)*. Springer-Verlag, Berlin, Heidelberg, 121–136. DOI: http://dx.doi.org/10.1007/978-3-642-25920-3_9
- Philip Levis, Sam Madden, Joseph Polastre, Robert Szewczyk, Cameron Whitehouse, Alec Woo, David Gay, Jason Hill, Matt Welsh, Eric Brewer, and David Culler. 2005. TinyOS: An Operating System for Sensor Networks. In *Ambient Intelligence*, Werner Weber, JanM. Rabaey, and Emile Aarts (Eds.). Springer Berlin Heidelberg, 115–148. DOI: http://dx.doi.org/10.1007/3-540-27139-2_7
- LinkSprite. 2009. JPEG Color Camera with Serial UART Interface. (2009). Retrieved Jul 27, 2013 from <http://www.linksprite.com/product/showproduct.php?lang=en&id=50>.
- Shaobo Liu, Jun Lu, Qing Wu, and Qinru Qiu. 2012. Harvesting-Aware Power Management for Real-Time Systems With Renewable Energy. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 20, 8 (August 2012), 1473–1486. DOI: <http://dx.doi.org/10.1109/TVLSI.2011.2159820>
- Jun Lu, Shaobo Liu, Qing Wu, and Qinru Qiu. 2010. Accurate modeling and prediction of energy availability in energy harvesting real-time embedded systems. In *Proceedings on the 1st Green Computing Conference (IEEE IGCC 2010)*. Chicago, IL, USA, 469–476. DOI: <http://dx.doi.org/10.1109/GREENCOMP.2010.5598280>
- Michele Magno, Stevan Marinkovic, Davide Brunelli, Emanuel Popovici, Brendan O’Flynn, and Luca Benini. 2012. Smart power unit with ultra low power radio trigger capabilities for wireless sensor networks. In *Proceedings of the Design, Automation and Test in Europe Conference (DATE 2012)*. Dresden, Germany, 75–80. DOI: <http://dx.doi.org/10.1109/DATE.2012.6176436>
- Maxwell Technologies, Inc. 2009a. Datasheet HC Power Series Ultracapacitors. (August 2009).
- Maxwell Technologies, Inc. 2009b. Design considerations for ultracapacitors, White paper. (April 2009).
- Clemens Moser, Lothar Thiele, Davide Brunelli, and Luca Benini. 2007a. Adaptive power management in energy harvesting systems. In *Proceedings of the Design, Automation and Test in Europe Conference (DATE 2007)*. Nice, France, 773–778. DOI: <http://dx.doi.org/10.1109/DATE.2007.364689>
- Clemens Moser, Lothar Thiele, Davide Brunelli, and Luca Benini. 2007b. Adaptive power management in energy harvesting systems. In *Proceedings of the Design, Automation and Test in Europe Conference (DATE 2007)*. EDA Consortium, San Jose, CA, USA, 773–778. DOI: <http://dx.doi.org/10.1109/DATE.2007.364689>
- David Moss, Jonathan Hui, Philip Levis, and Jung Il Choi. 2007. TEP 126: CC2420 Radio Stack. (June 2007). Retrieved Jul 27, 2013 from <http://www.tinyos.net/tinyos-2.x/doc/html/tep126.html>.
- Dong Kun Noh and Kyungtae Kang. 2011. Balanced energy allocation scheme for a solar-powered sensor system and its effects on network-wide performance. *J. Comput. System Sci.* 77, 5 (September 2011), 917–932. DOI: <http://dx.doi.org/10.1016/j.jcss.2010.08.008>
- Panasonic. 2008. Gold Capacitor HW Series Datasheet. (January 2008). Retrieved Jul 27, 2013 from <http://www.panasonic.com/industrial/components/pdf/ABC0000CE8.pdf>.
- Chulsung Park and Pai H. Chou. 2006a. AmbiMax: Autonomous Energy Harvesting Platform for Multi-Supply Wireless Sensor Nodes. In *IEEE SECON 2006*, Vol. 1. Reston, VA, USA, 168–177. DOI: <http://dx.doi.org/10.1109/SAHCN.2006.288421>
- Chulsung Park and Pai H. Chou. 2006b. A wearable wireless sensor platform for interactive art performance. In *In Proceedings of the Fourth Annual IEEE International Conference on Pervasive Computing and Communications (PerCom 2006)*. Pisa, Italy, 13–17. DOI: <http://dx.doi.org/10.1109/PERCOM.2006.12>

- Chulsung Park, Jinfeng Liu, and Pai H. Chou. 2005. Eco: an ultra-compact low-power wireless sensor node for real-time motion monitoring. In *Proceedings of the 4th international symposium on Information processing in sensor networks (IPSN 2005)*. Piscataway, NJ, USA. DOI: <http://dx.doi.org/10.1109/IPSN.2005.1440956>
- Diego Pizzocaro, Matthew P. Johnson, Hosam Rowaihy, Stuart Chalmers, Alun Preece, Amotz Bar-Noy, and Thomas La Porta. 2008. A knapsack approach to sensor-mission assignment with uncertain demands. In *Proceedings of SPIE, Unmanned/Unattended Sensors and Sensor Networks (SPIE 2008)*, Vol. 7112. Marseille, France. DOI: <http://dx.doi.org/doi:10.1117/12.799859>
- Alun Preece, Diego Pizzocaro, Konrad Borowiecki, Geeth de Mel, Mario Gomez, Wamberto Vasconcelos, Amotz Bar-Noy, Matthew P. Johnson, Thomas La Porta, Hosam Rowaihy, Gavin Pearson, and Tien Pham. 2008. Reasoning and resource allocation for sensor-mission assignment in a coalition context. In *Proceedings of the 2008 IEEE Military Communications Conference (IEEE MILCOM 2008)*. San Diego, California, 1–7. DOI: <http://dx.doi.org/10.1109/MILCOM.2008.4753114>
- Vijay Raghunathan, Aman Kansal, Jason Hsu, Jonathan Friedman, and Mani Srivastava. 2005. Design considerations for solar energy harvesting wireless embedded systems. In *Proceedings of the 4th international symposium on Information processing in sensor networks (IPSN 2005)*. Piscataway, NJ, USA. DOI: <http://dx.doi.org/10.1109/IPSN.2005.1440973>
- Ravishankar Rao, Sarma Vrudhula, and Daler N. Rakhmatov. 2003. Battery modeling for energy aware system design. *Computer* 36, 12 (Dec. 2003), 77–87. DOI: <http://dx.doi.org/10.1109/MC.2003.1250886>
- Aruna Ravinagarajan, Denis Dondi, and Tajana S. Rosing. 2010. DVFS based task scheduling in a harvesting WSN for structural health monitoring. In *Proceedings of the Design, Automation and Test in Europe Conference (DATE 2010)*. European Design and Automation Association, Dresden, Germany, 1518–1523. DOI: <http://dx.doi.org/10.1109/DATE.2010.5457052>
- Joaquin Recas Piorno, Carlo Bergonzini, David Atienza, and Tajana Simunic Rosing. 2009. Prediction and management in energy harvested wireless sensor nodes. In *The first International Conference on Wireless Communications, Vehicular Technology, Information Theory and Aerospace & Electronic Systems Technology (CTIF Wireless Vitae 2009)*. Aalborg, Denmark, 6–10. DOI: <http://dx.doi.org/10.1109/WIRELESSVITAE.2009.5172412>
- Joaquin Recas Piorno, Carlo Bergonzini, David Atienza Alonso, and Tajana S. Rosing. 2010. HOLLOWS: A Power-Aware Task Scheduler for Energy Harvesting Sensor Nodes. *Journal of Intelligent Material Systems and Structures* 21, 12 (2010), 1300–1332. DOI: <http://dx.doi.org/10.1177/1045389X10377033>
- Christian Renner, Jrgen Jessen, and Volker Turau. 2009. Lifetime Prediction for Supercapacitor-powered Wireless Sensor Nodes. In *FGSN 2009*. Hamburg, Germany, 55–58.
- Hosam Rowaihy. 2012. Location Privacy and Energy Preservation in Sensor Allocation Systems. *International Journal of Distributed Sensor Networks* 2012 (2012), 10. DOI: <http://dx.doi.org/10.1155/2012/197592> Article ID 197592.
- Hosam Rowaihy, Sharanya Eswaran, Matthew P. Johnson, Dinesh Verma, Amotz Bar-Noy, Theodore Brown, and Thomas La Porta. 2007. A survey of sensor selection schemes in wireless sensor networks. In *the Defense and Security Symposium on Unattended Ground, Sea, and Air Sensor Technologies and Applications (DSS 2007)*. Orlando, FL, USA.
- Hosam Rowaihy, Matthew P. Johnson, Amotz Bar-Noy, Theodore Brown, and Thomas La Porta. 2008. Assigning Sensors to Competing Missions. In *Proceedings of the IEEE Global Telecommunications Conference (IEEE GLOBECOM 2008)*. New Orleans, LA, USA, 44–49. DOI: <http://dx.doi.org/doi:10.1109/GLOCOM.2008.ECP.17>
- Hosam Rowaihy, Matthew P. Johnson, Ou Liu, Amotz Bar-Noy, Theodore Brown, and Thomas La Porta. 2010. Sensor-mission assignment in wireless sensor networks. *ACM Transactions on Sensor Networks* 6 (July 2010), 36:1–36:33. DOI: <http://dx.doi.org/10.1109/SAHCN.2011.5984925>
- Hosam Rowaihy, Matthew P. Johnson, Diego Pizzocaro, Amotz Bar-Noy, Lance Kaplan, Thomas La Porta, and Alun Preece. 2009. Detection and Localization Sensor Assignment with Exact and Fuzzy Locations. In *Proceedings of the 5th IEEE International Conference on Distributed Computing in Sensor Systems (IEEE DCOSS 2009)*. Springer-Verlag, Marina Del Rey, California, 28–43. DOI: http://dx.doi.org/doi:10.1007/978-3-642-02085-8_3
- Kuei-Ping Shih, Yen-Da Chen, Chun-Wei Chiang, and Bo-Jun Liu. 2006. A Distributed Active Sensor Selection Scheme for Wireless Sensor Networks. In *Proceedings of the 11th IEEE Symposium on Computers and Communications (IEEE ISCC 2006)*. Washington, DC, USA, 923–928. DOI: <http://dx.doi.org/10.1109/ISCC.2006.7>
- Farhan Simjee and Pai H. Chou. 2006. Everlast: long-life, supercapacitor-operated wireless sensor node. In *Proceedings of the 11th ACM/IEEE International Symposium on Low Power Electronics and Design (ACM/IEEE ISLPED 2006)*. New York, NY, USA, 197–202. DOI: <http://dx.doi.org/10.1109/LPE.2006.4271835>

- Jamie Bradley Steck and Tajana Simunic Rosing. 2009. Adapting task utility in externally triggered energy harvesting wireless sensing systems. In *Proceedings of the 6th international conference on Networked sensing systems (INSS 2009)*. Piscataway, NJ, USA, 16–23. DOI: <http://dx.doi.org/10.1109/INSS.2009.5409959>
- USCRN 2011. U.S. Climate Reference Network (USCRN). (2011). Retrieved Jul 27, 2013 from <http://www.ncdc.noaa.gov/crn/>.
- Christopher M. Vigorito, Deepak Ganesan, and Andrew G. Barto. 2007. Adaptive Control of Duty Cycling in Energy-Harvesting Wireless Sensor Networks. In *IEEE SECON 2007*. San Diego, California, USA, 21–30. DOI: <http://dx.doi.org/10.1109/SAHCN.2007.4292814>
- Ting Zhu, Abedelaziz Mohaisen, Yi Ping, and Don Towsley. 2012. DEOS: Dynamic energy-oriented scheduling for sustainable wireless sensor networks. In *Proceedings IEEE INFOCOM 2012*. Orlando, FL, 2363–2371. DOI: <http://dx.doi.org/10.1109/INFCOM.2012.6195625>
- Ting Zhu, Ziguo Zhong, Tian He, and Zhi-Li Zhang. 2013. Energy-synchronized computing for sustainable sensor networks. *Ad Hoc Networks* 11, 4 (2013), 1392 – 1404. DOI: <http://dx.doi.org/10.1016/j.adhoc.2010.11.005>

Received XXX; revised YYY; accepted ZZZ