Available online at www.sciencedirect.com
ScienceDirect
Computer Networks 50 (2006) 3701–3720
Computer Networks
www.elsevier.com/locate/comnet

An efficient heuristic for selecting active nodes in wireless sensor networks

Flávia Delicato [a,b,d,*], Fábio Protti [a,c,1], Luci Pirmez [a], José Ferreira de Rezende [b,d,1]

[a] Núcleo de Computação Eletrônica, Universidade Federal do Rio de Janeiro, Caixa Postal 2324, CEP 20001-970 Rio de Janeiro, RJ, Brazil
[b] Grupo de Teleinformática e Automação, Universidade Federal do Rio de Janeiro, Caixa Postal 2324, CEP 20001-970 Rio de Janeiro, RJ, Brazil
[c] Instituto de Matemática, Universidade Federal do Rio de Janeiro, Caixa Postal 2324, CEP 20001-970 Rio de Janeiro, RJ, Brazil
[d] Programa de Engenharia Elétrica – COPPE, Universidade Federal do Rio de Janeiro, Caixa Postal 2324, CEP 20001-970 Rio de Janeiro, RJ, Brazil

Received 16 August 2005; received in revised form 17 January 2006; accepted 24 April 2006
Available online 17 May 2006

Responsible Editor: I.F. Akyilidiz

Abstract

Energy saving is a paramount concern in wireless sensor networks (WSNs). A strategy for energy saving is to cleverly manage the duty cycle of sensors, by dynamically activating different sets of sensors while non-active nodes are kept in a power save mode. We propose a simple and efficient approach for selecting active nodes in WSNs. Our primary goal is to maximize residual energy and application relevance of selected nodes to extend the network lifetime while meeting application-specific QoS requirements. We formalize the problem of node selection as a knapsack problem and adopt a greedy heuristic for solving it. An environmental monitoring application is chosen to derive some specific requirements. Analyses and simulations were performed and the impact of various parameters on the process of node selection was investigated. Results show that our approach outperforms a naïve scheme for node selection, achieving large energy savings while preserving QoS requirements.

Keywords: Management of wireless sensor networks; Knapsack problem

* Corresponding author. Present address: Departamento de Informática e Matemática Aplicada (DIMAp), Universidade Federal do Rio Grande do Norte, Campus Universitário, Lagoa Nova, 59082-970 Natal, RN, Brazil. Tel.: +55 84 3215 3814; fax: +55 84 32153813.
E-mail addresses: fdelicato@nce.ufrj.br (F. Delicato), fabiop@nce.ufrj.br (F. Protti), luci@nce.ufrj.br (L. Pirmez), rezende@gta.ufrj.br (J. Ferreira de Rezende).
[1] Partially supported by CNPq, Brazil.

1389-1286/$ - see front matter
doi:10.1016/j.comnet.2006.04.005

1. Introduction

The design and deployment of wireless sensor networks (WSNs) constitutes a new domain of distributed computing that attracted great research interest in the last few years. In general, a WSN consists of a large number of low cost and densely deployed battery powered sensor nodes endowed with wireless communication, sensing, processing and storage capabilities [1]. WSNs are expected to be self-organized networks so that little or no setup is required. Their unique features make WSNs a powerful tool for environmental monitoring since they are able to perform geographically and temporally distributed in situ measurements [2].

Beyond environmental monitoring, there is a large range of applications for such networks that include: home automation [3], smart buildings [4], health and medical applications [5], vehicle and target tracking [6], among others. The majority of these applications requires a long network operational lifetime. For instance, applications of habitat monitoring may demand continuous operation through several months, and monitoring of civil structures (as bridges) may require an operational lifetime of several years [7]. However, the finite and non-renewable battery power of sensor nodes severely constrain the network lifetime [1].

Nonetheless, an important design feature of WSNs that can be useful for extending network lifetime is their high density of nodes. On one hand, this feature increases robustness against node failures and decreases energy costs by favoring multiple hop communications. On the other hand, this very same characteristic leads to a large redundancy in sensor-collected data. Recent researches [7–9] exploit this latter aspect by dynamically selecting a reduced set of sensors to remain active in the execution of a sensing task.

In this paper, we evaluate the benefits of adopting an enhanced mechanism for selecting nodes to be active as long as they fulfill application requirements. In other words, application-specific QoS requirements are considered when choosing nodes to be activated. The selection process is formulated as a knapsack problem [10], which is solved by a simple and efficient greedy heuristic. The main goal of the proposed solution is to maximize the relevance of the selected nodes (from the application point of view) and their residual energy, constrained by connectivity, coverage and energy issues [7].

The application class greatly influences how system resources have to be allocated to meet the requested level of QoS. For instance, applications of environmental monitoring [2], focus of this work, do not impose strict restrictions on data delay, because they are not time-critical but, on the other hand, they often require high data accuracy and long network lifetime.

Several works address the problem of active node selection, also known as node scheduling, in WSNs [8,9]. However, at the best of our knowledge, our work is the first attempt in the sense of encompassing application-specific requirements in the process of node selection. Our proposal has a low computational complexity, which makes it suitable for running on constrained sensor devices.

The remainder of this paper is organized as follows. In Section 2 we present the related works. In Section 3 we present the problem description and formulation. Section 4 describes the several assumptions and models adopted in the performed simulations. Section 5 describes in details the sets of simulations and the analysis of the obtained numerical results. Section 6 discusses the implications of a distributed approach for the selection of active nodes, in contrast with the presented centralized approach. Finally, Section 7 concludes the paper.

2. Related works

In the last years, several researchers have been investigating the problem of active node selection in WSNs [8,9,11,12], most of them with the aim of achieving high levels of energy efficiency. The majority of works considers coverage and connectivity guarantees as the unique QoS requirements for WSNs. In fact, the problem of sensing coverage has been extensively investigated. Several algorithms were proposed to find near to optimal solutions based on global information. In [12,13], techniques of linear programming are used to select the minimum set of active nodes able to maintain the complete sensing coverage of the network. Other protocols, such as GAF [14], AFECA [15] and ASCENT [16] aim to guarantee the network connectivity, but they do not address sensing coverage. SPAN [17] allows that sensors may be turned off whenever they are not being used as data sources or playing vital part in data routing. In [7], a solution is provided to meet both coverage and connectivity requirements. However, none of those protocols seeks for a balance between the quality of data generated for the application and the energy consumption of the network.

In [8,9] the problem of maximizing the lifetime of a WSN while guaranteeing a minimum level of quality at the application level is addressed. In those works, the problems of node selection and data routing are jointly addressed, and solved as a problem of generalized maximum flow. Those works present both an optimal and a heuristic solution with a totally centralized approach, based on global information. In contrast with the works discussed above, our work assumes that the scheme of active node selection is independent from the network routing protocol. Differently of approaches based on computational intensive techniques of linear programming, which are restricted to run off-line, our approach is light enough to be executed on-line and inside of the network.

3. Problem description and formulation

Given an application submitting a sensing task T to a WSN, the node selection process corresponds to the algorithm that decides which sensors should be active for the execution of that particular task. In order to avoid an early energy depletion of active nodes, the algorithm should alternate between subsets of active nodes during the complete task execution. In our algorithm, execution time is divided into rounds of size t. During each round the subset of selected nodes and their roles do not change. The decisions made by our algorithm are based upon information contained in interests submitted by the application, which consist of the task descriptor and QoS requirements. The former contains the type of sensor-collected data, the data-sending rate, the geographical area of interest (target area), and the monitoring duration and interval. The QoS requirements are application-dependent but, in the case of environmental monitoring, they may be expressed as minimum values for accuracy and spatial precision of sensor-collected data.

The node selection algorithm is executed in the following three cases: (i) initially, when a new application submits its interests to the network; (ii) proactively, for purposes of energy saving or due to changes in the application QoS requirements; or (iii) reactively, whenever some QoS violation is detected. The proposed scheme is based on the knapsack optimization problem [10] and it aims to maximize the lifetime of a network while guaranteeing a certain level of QoS to the application. The algorithm seeks to select the best subset of sensors to be activated by using three strategies: (i) minimizing network energy consumption by choosing the smallest possible number of nodes capable of providing the requested level of QoS; (ii) maximizing the sum of the residual energy of selected nodes, so that energy is spent in a uniform way among sensors during task execution time, thus avoiding the premature collapse of excessively used nodes; and (iii) taking into account the potential relevance, from the application point of view, of each individual sensor node.

The knapsack problem [10] can be stated as follows: given a non-negative real number M and a set of N objects where each object $i \in \{1,\dots,N\}$ is assigned a pair v_i, w_i of non-negative real numbers, we wish to find a subset $S \subseteq \{1,\dots,N\}$ that maximizes $v(S)$ subject to the constraint $w(S) \leqslant M$, where $v(S) = \sum_{i\in S} v_i$ and $w(S) = \sum_{i\in S} w_i$. The numbers v_i and w_i can be interpreted, respectively, as the utility and the weight of an object i. The number M can be interpreted as the capacity of a knapsack, that is, the maximum weight the knapsack can hold. Our goal consists of finding a collection of objects, the most valuable possible, which respects the capacity of the knapsack. In other words, the knapsack algorithm maximizes the utility of the objects placed in a knapsack of limited weight. It is well known that the knapsack problem can be solved by a dynamic programming algorithm in pseudo-polynomial time [10].

By formulating the problem of active node selection as a knapsack problem, sensor nodes are the objects to be placed in the knapsack. The sum of their utilities is optimized under the constraint of a certain energy budget M, which consists of the total amount of energy allocated to the task execution. Therefore, the budget M represents the knapsack capacity. Since the knapsack capacity is defined as the amount of energy necessary to accomplish a specific task, the value w_i, or weight of node i, is also given in terms of energy. Therefore, w_i is the cost of energy (sensing and communication) of a node i whenever it participates in a sensing task T. Such cost depends, among other factors, on: the type of the sensing device; the operating mode of sensor i during the task execution; and the data acquisition rate (sensing and transmission rates).

The utility of a given node i is given by both its *potential relevance* R_i and its *residual energy* U_i. The knapsack-based selection algorithm seeks to maximize the relevance R_i and the final residual energy U_i of selected nodes. Eq. (1) gives the objective function of the problem

$$\begin{aligned}&\text{maximize } \sum x_i(\alpha R_i + \beta(U_i - w_i))\\&\text{subject to } \sum x_i w_i \leqslant M,\end{aligned} \tag{1}$$

where: $x_i = 1$ if sensor i is selected to participate in task T, and 0 otherwise; R_i is a variable indicating the relevance of a node i; w_i is the energy spent in the task by sensor i; $U_i - w_i$ denotes the final energy of sensor i, if it participates in T; α and β are coefficients used to balance the priorities given for each term of the equation, and they depend on the application QoS requirements.

The relevance R_i of a node is a parameter related to the class of the target application. For time-critical applications, for instance, relevance is tied to features of a node that favor its capability of delivering data with low delay. For applications of environmental monitoring, focus of our work, relevance is related to factors that contribute for the accuracy of data generated by the node. A data accuracy related relevance of a node depends on its physical and topological features, given by its nominal precision (P_i), the environmental noise of its measurements (F_i), the number of neighboring nodes (N_i) and its proximity to the target area (A_i). Each parameter contributes to R_i computation, with a different weight. The value of P_i is a physical feature of each sensor and has the smallest weight among all terms. The parameter F_i is mainly influenced by physical characteristics of the location where the sensor node is deployed. The parameter F_i is a normalized value that depends on the actual level of environmental noise S_i, where S_i ranges from 0 to 100. The measurement environmental noise is then given by the following equation:

$$F_i = 1 - S_i/100. \tag{2}$$

The largest weights are assigned to parameters A_i and N_i. The value of N_i is inversely proportional to the amount of neighbors of the sensor, that is, it is proportional to the contribution of that node to sensing such location. To calculate A_i, sensors with distances d_i from the target area larger than the sensing range SR are automatically considered non-eligible. In order to assign a smaller value of relevance to sensors located at larger distances from the target area, we applied the formula:

$$A_i = 1 - d_i/\mathrm{SR}. \tag{3}$$

The values of parameters A_i and N_i are highly correlated. For instance, a sensor very close to the target area which has a small number of neighbors is likely to have a high relevance for the sensing task. On the other hand, a sensor very far from the target area and with a high number of neighbors probably has a smaller relevance from the application point of view. Considering the different weights of each parameter in the calculation of R_i and the correlation between A_i and N_i, the following equation is used:

$$R_i = \delta P_i + \phi F_i + \gamma/(A_i N_i), \tag{4}$$

where δ, ϕ and γ are coefficients representing the weights, $\delta < \phi < \gamma$.

The choice of active nodes in a WSN is subject to a set of constraints, described below, which should hold in any node selection scheme.

3.1. Energy constraints

A first constraint (R_1) is related to the fact that the amount of energy available in the network is a finite resource. At each round r, the energy spent by the selected subset of sensors cannot be larger than the network energy budget for such a round

$$\text{for all } r, \quad \sum_{i=1}^{N} x_i\ w_i^r\ t_r \leqslant Q_r, \tag{R$_1$}$$

where: N is the total number of sensors in the network; x_i is 0 or 1, depending whether sensor i was selected ($x_i = 1$) or not ($x_i = 0$) to be active at round r; w_i^r is the power consumption of sensor i during round r; t_r is the duration of round r; and Q_r is the energy budget for round r.

The proposed algorithm, based on the knapsack problem, naturally fulfills such a constraint, since the knapsack capacity represents the energy budget, i.e. the fraction of the total energy that the network/application wishes to allocate to a certain task during a given period of time (given by the number of rounds).

A second energy-related constraint (R_2) considers that a sensor node is only eligible to remain active in a round r if it has energy enough to remain alive until the end of the round. To meet such a constraint, a threshold L was defined as the minimum residual energy a node must have to be considered eligible. The constraint (R_2) can be defined as follows:

$$x_i \leqslant U_i/L, \tag{R$_2$}$$

where, again, x_i is a binary variable. If the residual energy U_i of sensor i is smaller than the threshold

L, x_i will be set to 0 (the sensor cannot be selected). Otherwise, x_i is free to be set to 1 (that is, sensor i is eligible). The constraint (R_2) is applied before the selection algorithm.

3.2. Coverage and connectivity constraints

Since the goal of a WSN is to monitor some geographic area, it has to maintain a full sensing coverage respecting a certain spatial precision, even when it operates in power save mode. We assume that a point p is covered by a node i if the Euclidian distance between them is smaller than the sensing range of the nodes, denoted by SR. A convex area A is said to have a coverage degree k (that is, A is k-covered) if every point p inside A is covered by at least k nodes [7].

Besides, a successful node selection scheme must also provide satisfactory connectivity so that active nodes can report collected data to the application. We assume that any two nodes i and j can communicate to each other only if the Euclidian distance between them is smaller than the radio range RR of the nodes, i.e., $\mathrm{dist}(i,j) < \mathrm{RR}$.

The coverage and connectivity constraints can be formulated as follows. Given a convex area A and a coverage degree k specified by the application, the number of inactive nodes should be maximized subject to the following constraints:

- the subset of active nodes guarantees that A is k-covered, i.e., for every point p of A:

  $$\sum_{i \in A(p)} x_i \geqslant k, \qquad (\mathrm{R}_3)$$

  where, $A(p) = \{i | \mathrm{dist}(p, i) < \mathrm{SR}\}$.
- *Connectivity assurance*: all active nodes have a valid route towards the data destination node D, that is, for every active node i, there exists a sequence $j_1, j_2, \ldots, j_n$ of active nodes such that the condition $\mathrm{dist}(j_q, j_{q+1}) < \mathrm{RR}$ holds for all $q = 1 \ldots n-1$, where $j_1 = i$ and $j_n = D$ (R_4).

To satisfy such constraints, a two-step procedure based on the disk-covering algorithm [18] is employed before executing the knapsack-based algorithm. In the first step, the target area (a rectangular area defined by the application) is totally covered by disks whose diameters are defined as the spatial precision requested by the application. Afterwards, the procedure heuristically selects k nodes that must remain active inside each disk. That selection takes into account the residual energy of the nodes. In the second step, the sensor field is totally covered by disks whose radii are equal to the radio range RR. To assure network connectivity, the procedure should guarantee that within each disk there is at least one active node.

3.3. Including QoS profiles

The basic QoS requirement of an application of environmental monitoring is the accuracy of the data supplied by the WSN. However, for most of such applications, the network lifetime is also of paramount relevance, since often a long time of monitoring is needed to correctly capture the temporal variations of long life cycle phenomena. Therefore, the application can choose to prioritize lifetime in favor of accuracy, or the opposite, to prioritize accuracy in favor of monitoring period, or it can choose to balance both parameters. Thus, in our work, the application QoS requirements, along with the parameter that it chooses to prioritize, compose a QoS profile. There are three possible QoS profiles:

1. Precision-based—it prioritizes the data accuracy or precision.
2. Lifetime-based—it prioritizes the network lifetime.
3. Ratio-based—it balances the network lifetime and the supplied accuracy, that is, it seeks the best tradeoff between energy consumption and data accuracy.

Considering the above QoS profiles, the original objective function is modified to include different weights according to the priority given by the application to the different QoS parameters. For precision-based profiles, larger values are assigned to the coefficient α; for lifetime-based profiles, larger values are assigned to the coefficient β; and finally, for ratio-based profiles, equal values are assigned to both coefficients.

4. Simulation models

The performance benefits of the proposed active node selection scheme were evaluated through simulations performed with JIST [19], a Java-based discrete event simulator. The simulation scenario assumed an application of environmental monitoring, which asks for raw data collection of a given physical phenomenon, in a target area during a given period of time. Additionally, the requested sensing task must meet the following requirements: (i) a minimum spatial resolution of $40\,m^2$ with a 1-coverage degree (that is, at least one node for each area of $40\,m^2$); (ii) a data acquisition rate of one sampling at each 10 s; and (iii) a data accuracy above a predefined threshold. As we seen in Section 3, data accuracy is directly related to the parameter R_i, first term in the objective function adopted by the selection scheme (Eq. (1)), which reflects the potential relevance of a node i from the application point of view. In our model, data accuracy is measured by the mean square error (MSE) value, which is given by the difference between a set of values assumed as "real" values of the monitored phenomenon and the set of values generated by the active sensors. The WSN lifetime should be long enough to guarantee data acquisition during the whole requested period of time as long as the stated QoS is respected. The lifetime is directly related to the value of network residual energy U_i, second term in the Eq. (4). Such value depends on the residual energy of the nodes activated for a given task and on their energy consumption during the task. The energy consumption, in turns, reflects the weight w_i of the node and it depends on the adopted energy model.

In the following subsections the models that characterize the network, the application, and the physical phenomenon used in our simulations are described. Next, simulation results are presented and discussed.

4.1. Network model

A WSN is usually composed of hundreds of sensor nodes and one or more sink nodes. Sink nodes are devices not limited by energy constraints and are endowed with high processing power. They act as entry points for application requests and as gathering points of sensor-collected data.

Our network model considers that all nodes are aware of their geographical positions and of their neighbors. The position of a sensor node can be obtained through the use of GPS or triangulation algorithms [20]. The neighbors' positions can be piggybacked (i) on initial configuration messages at the network initialization; (ii) on requested hello messages; or (iii) on sensing data messages.

Two different channels are assumed: a communication channel and a paging channel. The paging channel has low bandwidth and it is used to implement "wake up" and temporal synchronization schemes [21] used by each sensor.

All nodes are equipped with radios with the same maximum transmission range. Besides, nodes use a power control scheme in such a way that they always transmit with the minimum power needed to reach the next hop.

The data communication is performed through multiple hops from the data source to the sink. Intermediary nodes perform data aggregation whenever required by the application. The area that each sensor is able to monitor (sensing range) is defined as the circular area around the sensor with radius equal to the sensor sensing range.

The energy model assumes that all the sensors are capable of operating in a sleep/inactive mode or according to a number of predefined active modes. Active modes considered in this work refer to the role a sensor plays for a given sensing task. Sensors may assume three different roles: (i) source-only, for nodes placed inside the target area which solely generate and send their own data; (ii) router-only for nodes outside the target area; (iii) source/router, for nodes inside the target area which both generate and forward data. In each mode, a sensor node spends a different amount of energy [22]. A sensor in the inactive mode consumes an insignificant amount of energy. The different roles and modes of a sensor node are represented in the Table 1.

4.2. Application model

An application of environmental monitoring, demanding continuous measurements about a given

Table 1
Sensor operating modes and roles

Mode	Role	Sensing device	Processor	Transmitter	Receiver
Inactive	–	Off	Off	Off	Off
Active	Source-only	On	On	On	Off
Active	Router-only	Off	On	On	On
Active	Source/router	On	On	On	On

physical phenomenon, was chosen as the target of our work. The application defines a data-sending rate, a geographical area of interest, the total monitoring time and, optionally, one or more aggregation functions to be applied over the collected raw data. Furthermore, the application defines a minimum value for the accuracy and for the spatial precision of the sensor-collected data. An example of a task submitted by this category of application would be: "report average, minimum and maximum values of temperature, for the next 24 h, by sampling the area at every 60 s, with a maximum error of 5% and a spatial resolution of 20 m^2".

4.3. Physical phenomenon model

In order to generate data measurements as close as possible to the real world, the physical phenomenon being monitored by the simulated WSN needs to be defined. However, in order not to restrict the scheme to specific scenarios, a generic model for data dissemination from physical processes was adopted. In this model, if there is a data source j located at some point, its value V_j is diffused in the environment according to a power of the distance [23]. Therefore, the actual values reported by the sensors are a function of their nominal precision, the distance to each data source and the environmental noise associated to the measurement. Assume that there exist s data sources spread all over the environment. The actual value V_i^{actual} reported by a sensor node i located at some point p is given by the following equation:

$$V_i^{\text{actual}} = \frac{1}{s}\sum_{j=1}^{s}((a\ \text{dist}(i,j)+1)^{-b}\ V_j\ P_i + F_i), \qquad (5)$$

where: the indices $j = 1,\ldots,s$ stand for the data sources; dist(i,j) is the Euclidean distance between sensor i and data source j; a and b are adjusting parameters; V_j denotes the value of the "real" data generated by source j; P_i is the nominal precision of sensor i; and F_i is the environmental noise associated to the measurements. For the performed simulations we used $a = 0.25$ and $b = 1$ [23].

The current version of the model considers, at a first moment, only the spatial distribution of the physical phenomenon, disregarding its temporal variation. In this case, the data generated by the model supply a snapshot view of the monitored physical phenomenon. Subsequently, the time dimension was included in the model, in order to more realistically describe natural phenomena. Thus, besides varying along the geographical area, data representing the monitored phenomenon also changed along the time.

"Real" data were generated from a Gaussian distribution with mean 100 and variance 10. The nominal precision of each node is a random value uniformly distributed between 95 and 100 [2]. The environmental noise of each measurement was generated from a zero-mean Gaussian distribution with variance 1.

5. Analysis of results

The influence of a number of parameters in different scenarios was analyzed in the simulations in order to validate and evaluate the proposed scheme. In the next subsections the simulations are described in details and the obtained results are discussed.

5.1. Simulation description

A sensor field was created with 300 nodes randomly distributed in a square area with 200 m × 200 m. Each node has a radio range of 40 m and a sensing range of 20 m. The energy dissipation model adopted by the radio circuitries is as described in [22]. In such a model, the sleep mode power dissipation is about 416.3 mW, the idle time power dissipation is 727.5 mW, the receive power dissipation is 751.6 mW, and the transmit power dissipation is 986.0 mW. The target area was defined as a rectangular region of 100 m × 100 m inside the sensor field. Sensors located inside the target area were randomly selected to report data to sink nodes in the requested acquisition rate. A single sink node was placed in the top right corner of the field. Since we were not interested in simulating any specific communication protocol, we assumed hypothetical protocols, which deliver generated data from sources to the sink through the shortest path (in terms of geographic distance). We initially assumed that all transmissions were carried out without data loss. Further we introduced a data loss model in the simulations, in order to evaluate its impact on the performance of the proposed scheme. Each simulation runs for 1000 s, divided in 10 or more rounds, at the end of which the network residual energy was obtained and the MSE was calculated. The error bars shown in all presented graphs represent a confidence interval of 95%.

5.2. Analysis of the proposed scheme for active node selection

In the first stage of simulation, we compared the results of selecting different percentages of active nodes in terms of network residual energy and data accuracy. Since there is a single application using the network, the whole set of nodes could be allocated to supply the best possible QoS. However, we aimed to show that, by activating only a subset of nodes, the requested QoS could be met, saving network resources for further tasks and applications.

The adopted greedy heuristic for solving the knapsack problem consists of a two-steps algorithm: (i) first, the items (nodes) are sorted according to their computed relevance and residual energy; (ii) second, the items are inserted in the knapsack in the inverse order given by step (i), until an item does not fit in the knapsack (its weight is larger than the available capacity of the knapsack). The algorithm is greedy because it makes local decisions, which seems to be the most promising at the moment, and once the decision is made, it is never reconsidered.

As we seen in Section 3, when adapting the knapsack problem to the active node selection, the knapsack capacity represents a given budget, which means the amount of energy that should be spent by the whole network in the execution of a task. The total residual energy in a WSN is a value continually variable along the time, which depends on the number of active nodes in each moment and on several energy-consuming activities accomplished by such nodes. During the performed simulations, instead of considering budgets in terms of the absolute value of energy that should be spent by the network (sum of the energy spent by all selected nodes, with all their power consuming components), we chose to adopt an approach based on percentage of active nodes. This choice was made for the sake of the simplicity in the measurements to be taken. Thus, the network energy budget is specified as the percentage of nodes activated at each round. As the knapsack capacity stands for the budget allocated to perform a certain task, it corresponds to the sum of the weights of the respective percentage of activated nodes. In the adopted greedy approach, the weights of all nodes are assumed to be homogeneous and equal to their energy in the beginning of each round. Thus, from now on the knapsack capacity will be expressed as the percentage of nodes activated at each round.

The percentage of active nodes varied from 30% to 100% in the simulations. The initial energy of all nodes is a value randomly chosen and uniformly distributed between 15 J and 20 J. Nodes selected as inactive are completely turned off. Energy costs for turning off and restarting radios are considered negligible. The monitoring time requested by the application corresponds to 10 rounds and the maximum tolerated MSE is 0.3. The values considered for coefficients δ, ϕ and γ (Eq. (4)) were, respectively, 1, 2 and 3.

Figs. 1 and 2 show the network residual energy and the normalized MSE at the end of each simulation round for different budgets, i.e. percentages of active nodes, respectively.

Results show that a gain of 1000% in the residual energy is obtained when only 30% of nodes are activated, in contrast with activating 100% of nodes. From the 8th round on the MSE starts increasing for all budgets. This is due to a large number of sensors running short of energy. Lifetime expiration of source nodes or nodes located in the path from sources to the sink prevents data delivery. Although the MSE increases up to the 9th round for all budgets, it is still below the threshold tolerated by the application. From this point on and for budgets smaller than 70% the MSE increases to a value above the desired threshold, meaning that the application QoS is not being met anymore. Since the requested monitoring time was 10 rounds, results show that with only 70% of nodes the application QoS is met, resulting in large energy savings (more than 200%). Additional savings could be achieved by using smaller budgets, but at the expense of not meeting application QoS requirements.

5.3. Adaptation policies

For budgets smaller than 70%, the achieved error exceeded the established threshold before the end of the requested monitoring time, in spite of a significant amount of nodes remaining alive in the network. Those evidences suggest that (i) executing the knapsack-based selection procedure only at the first round—where the residual energy of nodes is very similar—may perform poorly; and (ii) adopting a fixed budget for the network during all rounds of a task may lead to an inefficient usage of the WSN resources or to the non-fulfillment of the application

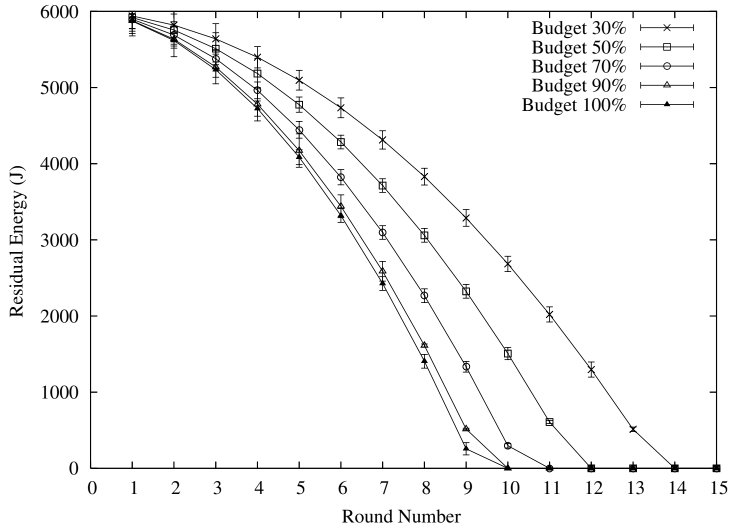

Fig. 1. Network residual energy at each round for each different budget.

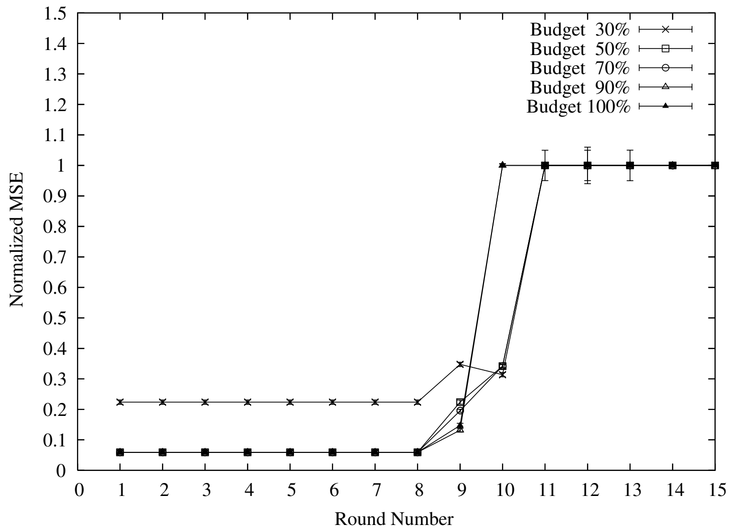

Fig. 2. Normalized MSE at each round for each different budget.

requirements, in spite of there being available resources in the network.

The strategy of increasing the initial budget whenever an error value is detected above the stipulated threshold could allow the networks to recover from QoS violations and meet the application needs for additional rounds. On the other hand, in cases where the error value is already below the threshold in the first rounds, additional energy savings could be obtained by decreasing the initial budget. Hence, we developed an adaptive strategy and evaluated its impact on the WSN performance.

In the developed strategy, an adaptation policy is implemented as a new run of the selection algorithm to be triggered on-demand, depending on parameters of the current network state. To reach such a goal, the MSE and the network residual energy values were monitored at the end of each round. If a threshold crossing is detected, the adaptation policy is triggered and the selection algorithm is executed either increasing or decreasing the budget. More specifically, whenever the error exceeds the threshold, the budget is increased by a factor of a. Also, if the MSE value is below the threshold, the budget is decreased by a factor of b. Several simulations were performed in order to analyze the effect of the coefficients a and b. The best results were obtained for $a = 0.3$ and $b = 0.1$.

Results of simulations performed for an initial budget of 100% show that without adaptation the WSN had its energy completely depleted at the end of the 13rd round. The MSE value was kept below the requested threshold up to the 10th round with both schemes. From this point on, the MSE dramatically grows beyond the desired threshold when no adaptation is adopted. The main reason for this is the lack of energy of source nodes, leading to less data being delivered to the sink node and the consequent error increase.

When adopting the adaptation policy, we observed that savings of up to 200% in the network final energy were obtained for initial budgets ranging from 100% to 70%, in comparison to the fixed budget strategy. Therefore, with the adaptation policy the network lifetime could be extended for additional rounds. For budgets smaller than 70% there were not significant advantages with the adaptive approach. The main reason for such behavior was that, with small budgets, the error values were already close to the requested threshold in the initial rounds, thus preventing a significant budget decrease in subsequent rounds. Concerning the error values, we observed that when the budget was increased in order to solve the QoS violation, the error did not necessarily decrease. The reason for such unexpected behavior was the selection of some nodes with high residual energy, in spite of their low accuracy. As an attempt for overcoming such a problem, we increased the priority of node data accuracy. To achieve this goal, we changed the value of coefficient α in the Eq. (1) to have a larger weight than the coefficient β. Recall that in the proposed scheme, the utility of a node is given by both its residual energy and the potential relevance of its data for the application. Hence, relevance is a measure directly related to the provided data accuracy. We also increased the weights of the terms that reflect the precision of the sensor node and the environmental noise (P_i and F_i in Eq. (4), respectively).

Results in Figs. 3 and 4 show that by applying these improvements in the selection process along with the adaptation policy, the WSN lifetime could be extended for up to 20 rounds, while the error is kept below the requested threshold at least until the 21st round. It is important to point out that at each new run of the selection algorithm triggered by the adaptation scheme, a different set of nodes were selected to act as routers as well as sources. These new selected source nodes, according to their nominal precision and distance to the phenomenon, provide different values for data accuracy. Such behavior leads to the oscillations in the MSE values that can be observed in the curve of Fig. 4.

In spite of these improvements in the node selection process, nodes with low accuracy were still selected due to their high residual energy. We observed that having high residual energy makes these nodes good choices as routers, but not as data

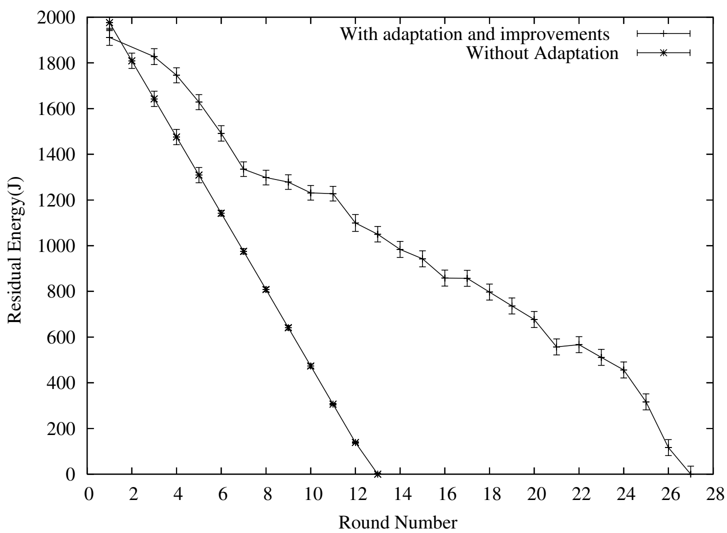

Fig. 3. Network residual energy at each round, without adaptation and with adaptation/improvements (initial budget 100%).

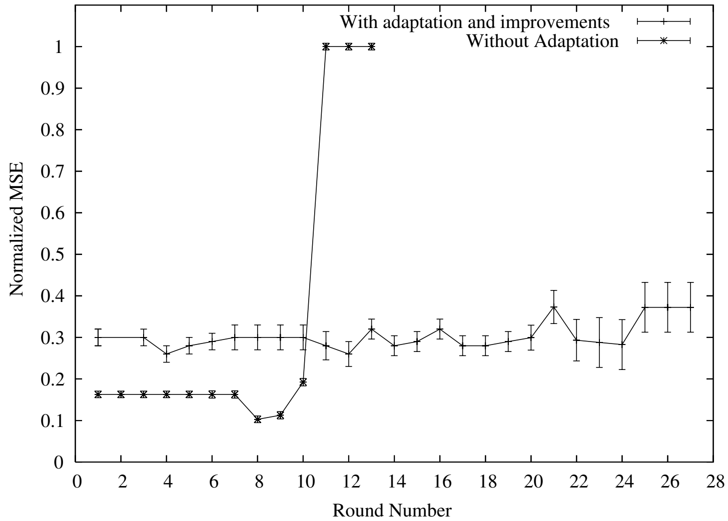

Fig. 4. Normalized MSE at each round, without adaptation and with adaptation/improvements (initial budget 100%).

sources. However, the selection algorithm does not distinguish among the roles an active node is assigned to. As a possible solution for this issue, we modified the algorithm to separate the roles of the nodes among router and sensor nodes during the selection process. This selection algorithm variant is executed in two steps, each one considering different criteria for choosing active nodes. In other words, two different knapsacks should be filled.

5.4. *Using two knapsacks in the selection process*

This stage of simulation aims to exploit the fact that to assume the role of source or router requires different characteristics of the nodes. Nodes with high potential to play the role of a router should have larger number of neighbors (considering the radio range) and high residual energy. On the other hand, nodes with great potential to act as sources should have a small number of neighbors (considering the sensing range) and provide high data accuracy. Therefore, the selection algorithm was split in two-steps, in which two different objective functions were used. In the first step, only nodes that will act as sources were selected. The objective function was changed to maximize the weight of the most important features for such a role. In the second step, the algorithm selects nodes that will act as routers according to an objective function that considers important attributes for this role. Nodes that were also selected in the first step will assume the router/source role.

Simulations were run using such variant of the selection algorithm along with the adaptive policies described in Section 5.3. Results are shown in Figs. 5 and 6. The network lifetime is extended to more than 30 rounds and up to the 28th round both the QoS and sensing coverage requirements are met. An unexpected behavior was observed quite often when the network budget is decreased as a strategy for saving energy. The MSE value decreased instead of increasing as occurred with the one-knapsack variant. Such behavior may be explained by the fact that the smaller is the budget used the better is the choice of sources according to relevant criteria to this role. In other words, when new nodes are included thanks to larger budgets, the resulting data accuracy degrades due to worse values provided by these nodes.

5.5. *QoS profiles*

All previous simulations assumed a ratio-based QoS profile. In such a profile, values of both coefficients α and β were set to 1. Next simulations evaluate the effect of using the different profiles described in Section 3.3. For the precision-based profile the value of the coefficient α was set to 50, while β was set to 1. For the lifetime-based profile the value of the coefficient α was set to 1, while β was set to 50. For the ratio-based profile the values of both the coefficients were set to 1.

Initially, simulations were run with the same configuration as the previous experiments. Recall

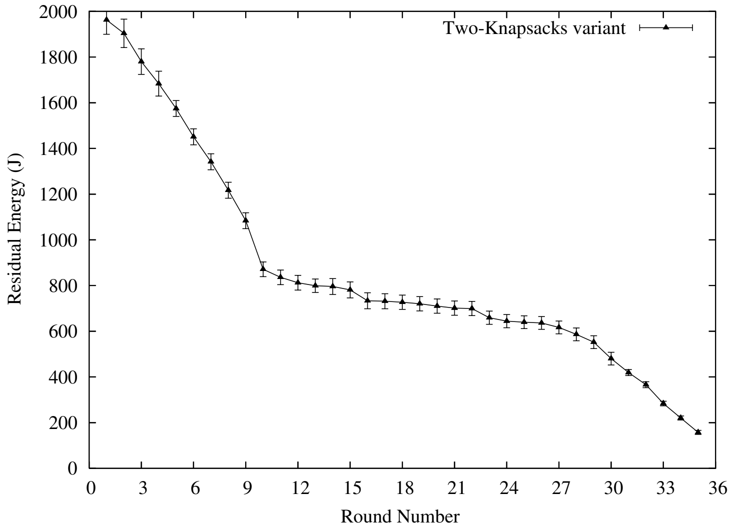

Fig. 5. Residual energy at the end of each round using the two-knapsacks variant.

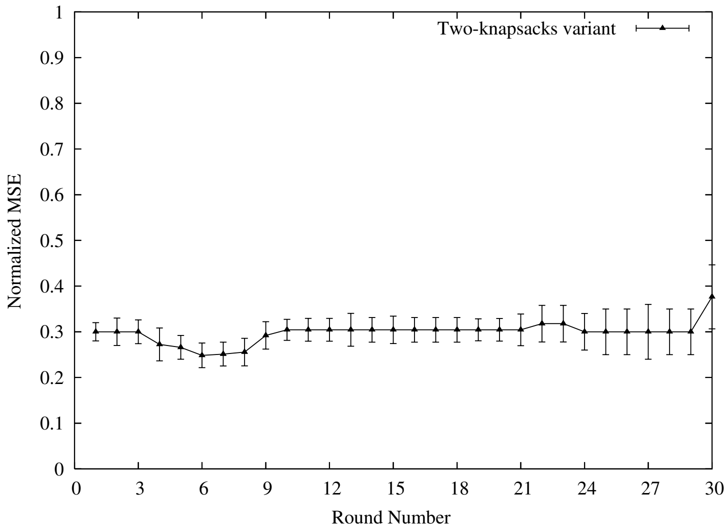

Fig. 6. Normalized MSE at the end of each round using the two-knapsacks variant.

that the initial energy of all nodes is a random value ranging from 15 J to 20 J. With such a configuration, the achieved results show that the final energy does not significantly change among the different profiles. We impute this result to the fact that the selection algorithm runs before the first round, when the residual energy of all nodes is very similar. In a subsequent set of simulations, we set the values of node initial energy as randomly chosen between 0 J and 20 J, in order to have more significant variation among such values.

Figs. 7 and 8 depict the residual energy and the MSE value at the end of the 10th round for each budget and each QoS profile, respectively, using this new parameter for node initial energy. Results show that for the lifetime-based profile there was savings of up to 50% in the network final energy in comparison with the precision-based profile. Regarding the MSE values, when the application decides to prioritize the relevance, which is directly related to the provided data accuracy (precision-based profile), the final value of error is significantly smaller than the one when the network lifetime is prioritized.

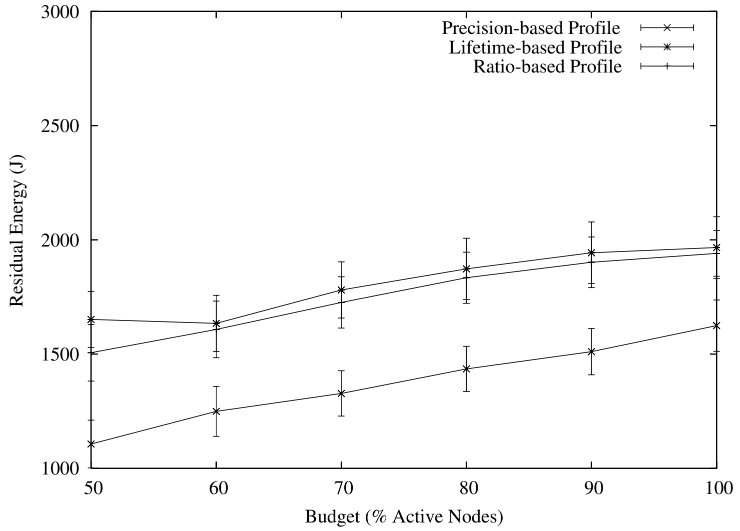

Fig. 7. Network residual energy at the 10th round for the different budgets and QoS profiles.

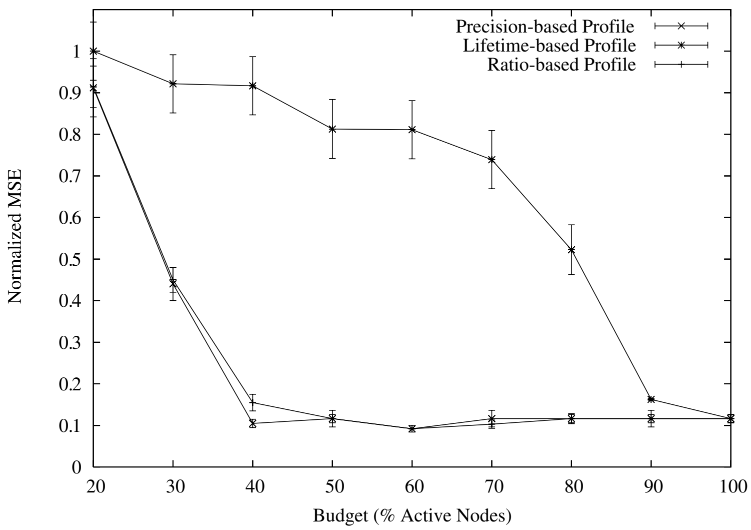

Fig. 8. Normalized MSE at the 10th round for the different budgets and QoS profiles.

5.6. Analyzing the impact of data losses

The proposed scheme for node selection is independent from the underlying communication protocols (of both the network and MAC layers) and most of the performed simulations do not consider the occurrence of losses in the data transmissions. However, in real environments of WSNs, there are several factors that may cause data loss. Data may be lost due to interferences or collisions at the MAC level. Furthermore, since sensor nodes are prone to malfunctioning, data may also be lost due to permanent or temporary failures of communicating nodes.

To study the impact of data losses on the proposed scheme, we introduced a random (uniform distribution) data loss model between each source node and the sink node. In this model, losses are uncorrelated and each data transmitted by the source is lost with a probability p. The probability associated to each source ranges from 0% to 50%. Fig. 9 depicts the values of MSE at the last round

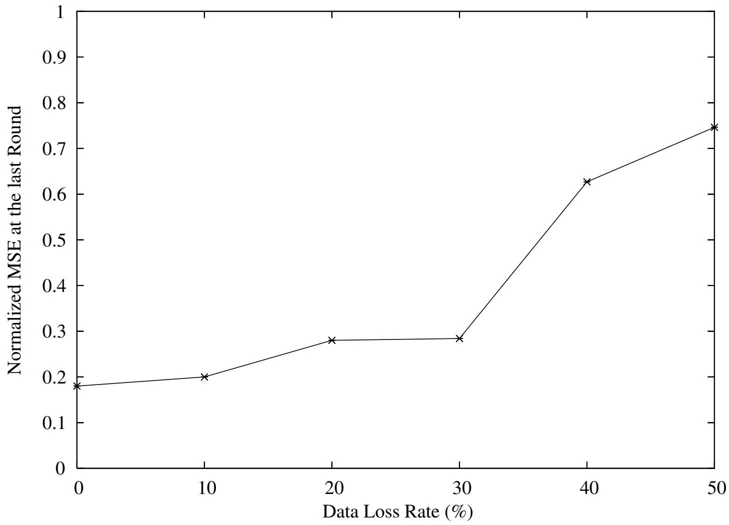

Fig. 9. Normalized MSE at the 10th round for the different percentages of data losses.

for each probability. Results show no significant difference among MSE values for 0–10% of data loss. This behavior is explained by the inherent redundancy in the data values sent by sensor nodes, which guarantees that the accuracy does not decrease. From 10% up to 30% of data losses the MSE value increased, consequently decreasing the data accuracy, which, however, remained below the threshold requested by the application. From 30% on, the error value rises above the threshold. Therefore, simulations demonstrated that the proposed scheme is robust to up to 30% of data loss, without compromising the final data accuracy.

5.7. Analyzing the temporal variation of the physical phenomenon

All the previous simulations considered only spatial variation of the data representing the monitored phenomenon. In this stage of simulation, the time dimension was included in order to analyze the behavior of the proposed scheme with more realistic scenarios, in which sensor-collected data is both spatial and time-variant. Simulations were performed as follows. In each single simulation, rounds were divided into time intervals of equal duration, and real data values associated to the physical phenomenon were changed at each time interval. New data values were generated at each interval according to the same Gaussian distribution described in Section 4.3. The duration of the time interval, which, in fact, represents the temporal variability of the phenomenon, was simulated ranging from 50 to 250 s. The impact of such variability was evaluated in terms of data accuracy values at the end of each simulation round. Results show that no significant difference occurred in accuracy values related to the temporal variability of the phenomenon. This behavior was as expected, since the final accuracy of the data delivered by the network in a given round is primarily related to which nodes were selected to be active in that round. In cases when data values are time variant, the collecting and sending rates have also influence on the data accuracy. Since such rates are often several orders of magnitude smaller than the temporal variability of whatever existent physical phenomenon, it is not likely that the interval of changing data values affects the generated MSE.

5.8. Analyzing the impact of round duration

As we seen in Section 3, the execution time of a sensing task is divided in rounds of duration t, during which the set of active nodes does not change. The value of t is an important parameter, which may have influence on the behavior of the proposed scheme. Values too small for t lead to the frequent reorganization of the logical topology, generating oscillations in the network and consuming energy due to the exchange of control messages. On the other hand, values too large for t may lead to the decrease of the number of nodes potentially eligible for being activated, since a node must have residual energy enough to remain alive during the entire round extent. Thus, there is a relevant tradeoff that should be analyzed when assigning the value for t. We evaluated such tradeoff by carrying out a set of simulations varying the value of t. The different values for the round duration were evaluated in terms of their impact over the final values of network residual energy and the generated MSE (which reflects the data accuracy). We also analyzed the round duration against the different values of the temporal variability of the physical phenomenon, in order to verify if there is any correlation between these two variables. Results are shown in Figs. 10 and 11.

The different values for the phenomenon temporal variability presented the same behavior, independent from the value assigned to t, demonstrating the inexistence of a correlation between these two para-meters, for the kind of simulated phenomena. Regarding the values of residual energy, results show that the network final energy progressively decreased with the increase of t. This is an obvious result, since the set of active nodes, their roles and operation modes, do not change along a given round. Thus, with larger values for t, the same nodes consume energy during more time, reducing the whole network residual energy. Regarding the data accuracy, for

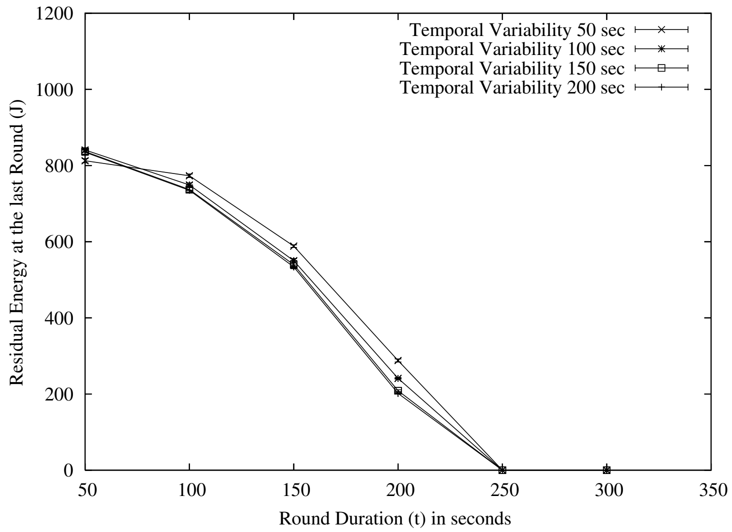

Fig. 10. Residual energy at the 10th round, for the different values for t and for the temporal variability of the physical phenomenon.

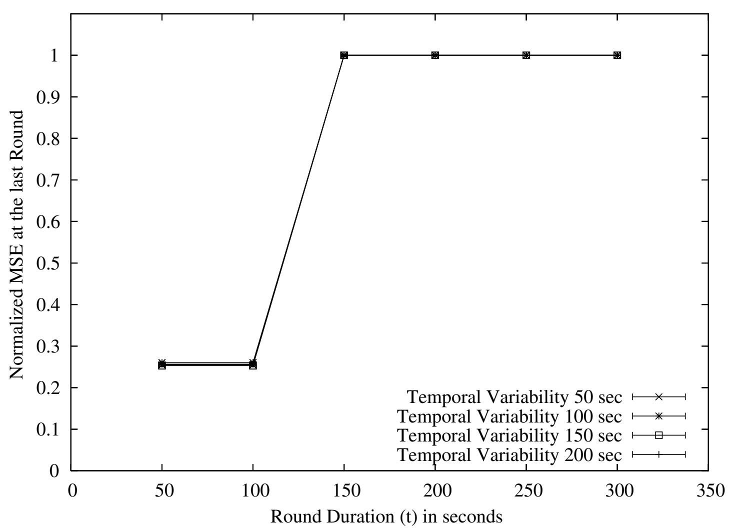

Fig. 11. Normalized MSE at the 10th round, for the different values for t and for the temporal variability of the physical phenomenon.

round durations from 50 to 100 s, there was no difference between the values of MSE. For values of t larger than 100 s, the MSE dramatically increased to values above the requested threshold. Such behavior occurred due to the restriction (R_2), which must hold in the proposed scheme. According to this restriction, a node is only eligible to be active in a round if it has residual energy enough to remain alive during the entire round. With larger duration, fewer nodes can be selected as active during the round. In spite of the inherent redundancy of the sensor-collected data, there is a boundary above which the amount of nodes reporting data affects negatively the generated error, consequently decreasing the provided accuracy.

5.9. Comparison with a Naïve approach for node selection

For the purpose of comparison, a naïve scheme for node selection was created as a baseline for our work. In such a scheme, for each budget (percentage of active nodes) the number of nodes corresponding to that percentage (in relation to the total number of nodes in the WSN) is allocated in a random fashion. Only the target area is taken into account, guaranteeing that the corresponding percentage of source nodes inside the target area is selected. However, neither the criterion of residual energy nor relevance-related criteria was considered in the process of selecting active nodes. Furthermore, the scheme did not offer any guarantees of sensing coverage or network connectivity. The goal of simulations performed with the naïve scheme was to evaluate the benefits of the strategy adopted by our work.

Figs. 12 and 13 show the residual energy and the normalized MSE at the end of 10 rounds for both the proposed algorithm and the naïve approach. It can be noted that for all budgets the final residual energy of the network was larger when the proposed scheme was adopted. The MSE was always smaller with the proposed approach than with the random selection of nodes. In the naïve approach, since the network connectivity is not assured, mainly with smaller budgets, source nodes could not deliver their data because they were not able to establish a route to the sink. The node selection without considering the node residual energy was also inefficient, since there was a smaller balance of energy consumption in the WSN, thus reducing its lifetime.

5.10. Analysis of tradeoffs accuracy-lifetime

Schemes for the intelligent node scheduling in WSNs, based on activating different percentages of nodes, such as the proposed in this work, allows setting network configurations that range from lots of active nodes/high accuracy and small lifetime/high energy consumption to few active nodes/low accuracy and long lifetime/low energy consumption, depending on the application-specific requirements. From simulation results, we will be able to pinpoint

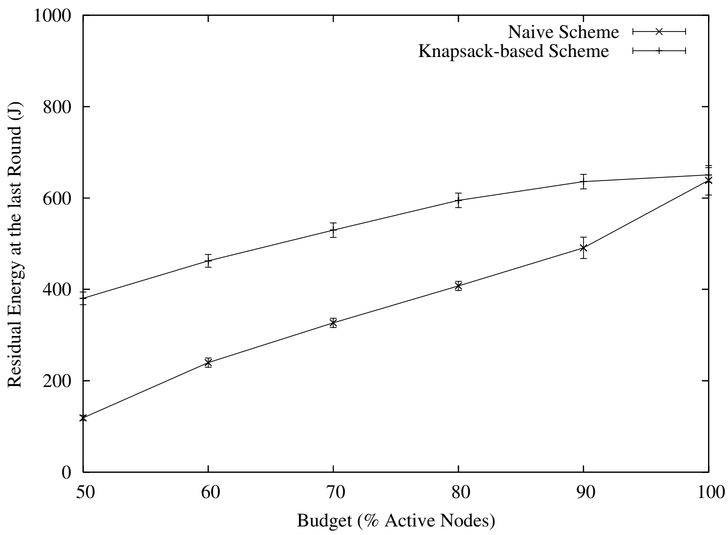

Fig. 12. Network residual energy at the 10th round, considering the proposed scheme for node selection and the naïve scheme.

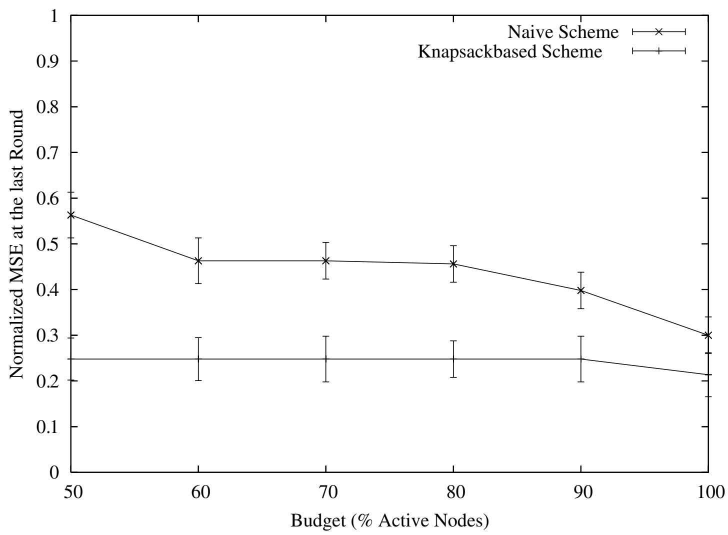

Fig. 13. Normalized MSE at the 10th round, considering the proposed scheme for node selection and the naïve scheme.

a number of tradeoffs among different requirements, which will assist network managers or application developers in their configuration decisions. Such tradeoffs could provide, for instance, the highest accuracy that can be obtained given an energy budget and a number of rounds. On the other hand, it is possible to establish a value for accuracy and the associated maximum expected network lifetime.

Fig. 14 provides an example of curve that could be examined by an application in order to know the configuration options the WSN offers, and the respective potential quality of results to be achieved in each configuration. By looking at the curve, the application can decide, for example, whether it will need to trigger adaptation policies to meet their QoS requirements during the monitoring task or not. Furthermore, if the application is interested in prioritizing data accuracy as its QoS requirement, the curve provides the approximate network lifetime (represented as number of rounds) that can be expected. By being aware of this estimated value, the application may choose relaxing the accuracy parameter or not. On the other hand, if the application prioritizes network lifetime, the curve provides the approximate maximum data accuracy expected to be achieved for each number of rounds.

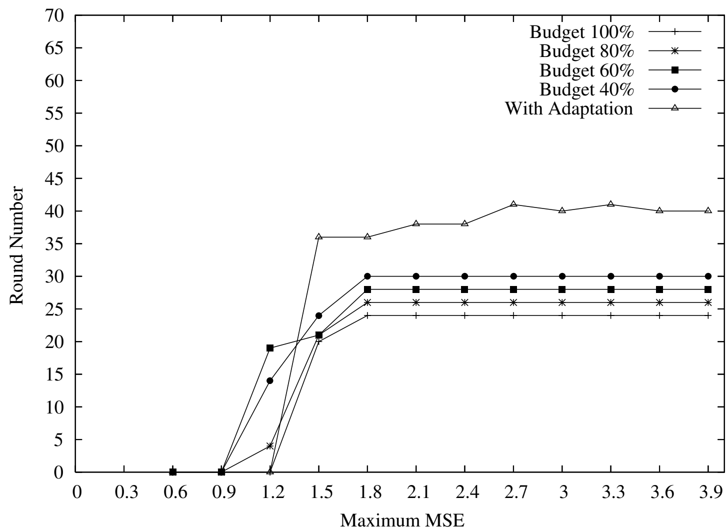

Fig. 14. Tradeoffs: network lifetime (in number of rounds) versus accuracy (in MSE values).

6. Towards a distributed approach

In the proposed scheme, the procedures for electing active nodes are performed in a centralized way; the decisions, however, do reflect different local behaviors in distinct parts of the whole WSN. In fact, a view of the potential relevancy, in terms of energy and quality of data, of different sets of geographically close groups of nodes can be achieved from the information sent by all sensors.

The simulations described in this article assume a flat topology and the decisions are made by sink nodes. However, if the network adopted a hierarchical topology, such decisions could be made by cluster-heads based on information of their cluster members only. This strategy does not incur any modification in the adopted algorithms, and has the advantage of improving scalability of the proposed scheme.

Nonetheless, it is well known that totally distributed solutions, based on localized algorithms, are the most suitable and robust options for WSNs, considering the common large-scale of such networks. Therefore, a distributed version of the proposed scheme for active node selection certainly would be a valuable contribution. In such version, the same parameters considered in the process of node selection with the centralized approach (nominal precision, residual energy, etc.) should be only exchanged among neighboring nodes. After receiving data from its neighborhood during a predefined time interval, each node would make the decision of whether or not to stay active upon through a local heuristic. Such heuristic would be derived from the same assumptions presented in this article, which aim to select nodes with more relevance for the application and higher residual energy. However, such a distributed solution is a totally new approach, requiring new algorithms and simulations. We intend to exploit this solution in a future work.

7. Conclusions

We proposed a simple and efficient approach for selecting active nodes in WSNs. Our primary goal was to maximize residual energy and application relevance of selected nodes to extend the network lifetime while meeting application-specific QoS requirements. We formalized the problem of node selection as a knapsack problem and we adopted a greedy heuristic for solving it. An application of monitoring environment was chosen to derive some specific requirements. We evaluated our proposal by performing several simulations and examining the impact of various parameters on the process of node selection. Results were promising, with large energy savings achieved while preserving QoS requirements. Furthermore, results show that our approach outperforms a naïve scheme for node selection, which does not take into account individual features of sensors.

The adoption of an enhanced mechanism for node scheduling in WSNs allows balancing the several tradeoffs between the QoS provided to WSN applications and the consumption of the scarce network resources. The task of node scheduling can be considered as a management service in WSNs. An intelligent node management service can be provided by WSN middleware systems to optimize the network operation, by directly interacting with the lowest levels of the protocol stack. Recent efforts have been accomplished in this track, and new middleware systems have been proposed to support such kind of optimization, as well as to simplify the process of application development [9,24–27]. We hope that the adoption of our scheme for node management as part of a WSN middleware can push the network lifetime up to its upper bounds, providing the best utilization of the network resources, by several different applications.

Future directions of our work comprise: (i) to create a new version of the knapsack-based algorithm which takes into account different QoS parameters for selecting nodes, such as data latency, and (ii) to implement a distributed version of the scheme for node selection, in which nodes take decisions regarding their activation based only on localized information sent by neighboring nodes.

References

[1] I. Akyildiz, W. Su, Y. Sankarasubramaniam, E. Cayirci, Wireless sensor networks: a survey, Computer Networks 39 (2002) 393–422.

[2] A. Mainwaring, J. Polastre, R. Szewczyk, D. Culler, J. Anderson, Wireless sensor network for habitat monitoring, in: Proc. of the ACM International Workshop on Wireless Sensor Networks and Applications, Atlanta, Georgia, USA, September 2002.

[3] E.M. Petriu, N.D. Georganas, D.C. Petriu, D. Makrakis, V.Z. Groza, Sensor-based information appliances, IEEE Instrumentation and Measurement Magazine 3 (2000) 31–35.

[4] G.D. Abowd, J.P.G. Sterbenz, Final report on the interagency workshop on research issues for smart environments, IEEE Personal Communications 1 (2000) 36–40.

[5] P. Bauer, M. Sichitiu, R. Istepanian, K. Prennaratne, The mobile patient: wireless distributed sensor networks for patient monitoring and care, in: Proc. of the IEEE International Conference on Information Technology Applications in Biomedicine, Washington DC, USA, November 2000.
[6] M. Duarte, Y.-H. Hu, Vehicle classification in distributed sensor networks, Journal of Parallel and Distributed Computing 64 (2004) 826–838.
[7] X. Wang, G. Xing, Y. Zhang, C. Lu, R. Pless, C. Gill, Integrated coverage and connectivity configuration in wireless sensor networks, in: Proc. of the ACM SenSys03, Los Angeles, California, USA, November 2003.
[8] M. Perillo, W. Heinzelman, Optimal sensor management under energy and reliability constraints, in: Proc. of the IEEE Wireless Communications and Networking Conference, New Orleans, Louisiana, USA, March 2003.
[9] M. Perillo, W. Heinzelman, Sensor management policies to provide application QoS, Elsevier AdHoc Networks Journal. 1 (2–3) (2003) 235–246.
[10] T.H. Cormen, C.E. Leiserson, R.L. Rivest, C. Stein, Introduction to Algorithms, MIT Press, 2001.
[11] J.Liu et al., A dual-space approach to tracking and sensor management in wireless sensor networks, in: Proc. of the ACM International Workshop on Wireless Sensor Networks and Applications, Atlanta, Georgia, USA, September 2002.
[12] S. Meguerdichian, M. Potkonjak, Low power 0/1 coverage and scheduling techniques in sensor networks, UCLA Technical Reports 030001, January 2003.
[13] K. Chakrabarty et al., Grid coverage for surveillance and target location in distributed sensor networks, IEEE Transactions on Computers 51 (12) (2002) 1448–1453.
[14] Y. Xu, J. Heidemann, D. Estrin, Geography-informed energy conservation for ad hoc routing, in: Proc. of the ACM SIGMOBILE Annual International Conference on Mobile Computing and Networking, Rome, Italy, July 2001.
[15] Y. Xu, J.Heidemann, D.Estrin, Adaptive energy-conserving routing for multihop ad hoc Networks, Research Report 527, USC/Information Sciences Institute, October 2000.
[16] A. Cerpa, D. Estrin, ASCENT: Adaptive Self-configuring Sensor Networks Topologies, in: Proc. of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies, New York, USA, June 2002.
[17] B. Chen et al., Span: an energy-efficient coordination algorithm for topology maintenance, in ad hoc wireless networks, in: Proc. of the ACM SIGMOBILE Annual International Conference on Mobile Computing and Networking, Rome, Italy, July 2001.
[18] J. Pach, P.K. Agarwal, Combinatorial Geometry, Wiley, New York, 1995.
[19] JIST—Java in Simulation Time. Available from: <http://jist.ece.cornell.edu/>.
[20] A. Nasipuri, K. Li, A directionality based location discovery scheme for wireless sensor networks, in: Proc. of the 1st ACM international workshop on Wireless sensor networks and applications, Atlanta, Georgia, USA, September 2002.
[21] J. Elson, L. Girod, D. Estrin, Fine-grained network time synchronization using reference broadcasts, in: Proc. of the 5th symposium on Operating systems design and implementation, Boston, Massachusetts, USA, December 2002.
[22] D. Estrin, A. Sayeed, M. Srivastava, Wireless sensor networks, Mobicom 2002 Tutorial. Available from: <http://www.sigmobile.org/mobicom/2002/program/tutorial.html>.
[23] A. Boulis, S. Ganeriwald, M. Srivastava, Aggregation in sensor networks: an energy-accuracy trade-off, in: Proc. of the First IEEE workshop on Sensor Network Protocols and Applications, Anchorage, Alaska, USA, May 2003.
[24] F.C. Delicato, P.F. Pires, L. Pirmez, L. Rust, A flexible middleware system for wireless sensor networks, in: Proc. of the ACM/IFIP/USENIX International Middleware Conference, Rio de Janeiro, Brazil, July 2003.
[25] F.C. Delicato, P.F. Pires, L. Pirmez, J.F. Rezende, Reflective middleware for wireless sensor networks, in: Proc. of the 20th ACM Symposium on Applied Computing, Santa Fe, New Mexico, USA, March 2005.
[26] K. Nahrstedt, D. Xu, QoS-aware middleware for ubiquitous and heterogeneous environments, IEEE Communications Magazine 39 (2001) 140–148.
[27] W. Heinzelman, A. Murphy, H. Carvalho, et al., Middleware to support sensor network applications, IEEE Network Magazine Special Issue 18 (2004) 6–14.

Flávia Coimbra Delicato received a M.Sc. degree on Computer Science in 2000 and a D.Sc. degree on Electrical Engineering in 2005 both from the Federal University of Rio de Janeiro (UFRJ), Brazil. She is currently an associate professor at Federal University of Rio Grande do Norte. Her research interests include wireless networks, wireless sensor networks, web services and middleware technologies.

Fábio Protti got his Ph.D. in 1998 at the Federal University of Rio de Janeiro (UFRJ), Brazil. Since then, he has taught at Federal University of Rio de Janeiro as an Associate Professor at the Computer Science Department. He is interested in graph theory and complexity of algorithms.

Luci Pirmez received a B.Sc degree on computer science in 1981, a M.Sc. degree on computer science in 1986 and the D.Sc. degree on computer science in 1996 from the Federal University of Rio de Janeiro, Brazil. She is a member of research staff of the Computer Center of Federal University of Rio de Janeiro. Her research interests include wireless networks, wireless sensor networks, networks management and security and formal description techniques for communication protocols.

José Ferreira de Rezende received the B.Sc. and M.Sc. degrees in Electronic Engineering from the Federal University of Rio de Janeiro (UFRJ), Brazil, in 1988 and 1991, respectively. He received the Ph.D. degree in Computer Science from the Université Pierre et Marie-Curie, Paris France, in 1997. He was an associate researcher at LIP6 (Laboratoire d'Informatique de Paris 6) during 1997. Since 1998 he is an Associate Professor at Federal University of Rio de Janeiro. His research interests are in wireless networks, distributed multimedia applications, multipeer communication, high speed and mobile networks and quality of service in the Internet.