Wireless Netw (2008) 14:831–858
DOI 10.1007/s11276-007-0017-x

Controlled sink mobility for prolonging wireless sensor networks lifetime

Stefano Basagni · Alessio Carosi · Emanuel Melachrinoudis · Chiara Petrioli · Z. Maria Wang

Published online: 7 February 2007

Abstract This paper demonstrates the advantages of using controlled mobility in wireless sensor networks (WSNs) for increasing their lifetime, i.e., the period of time the network is able to provide its intended functionalities. More specifically, for WSNs that comprise a large number of statically placed sensor nodes transmitting data to a collection point (the sink), we show that by controlling the sink movements we can obtain remarkable lifetime improvements. In order to determine sink movements, we first define a Mixed Integer Linear Programming (MILP) analytical model whose solution determines those sink routes that maximize network lifetime. Our contribution expands further by defining the first heuristics for controlled sink movements that are fully distributed and localized. Our *Greedy Maximum Residual Energy* (GMRE) heuristic moves the sink from its current location to a new site as if drawn toward the area where nodes have the highest residual energy. We also introduce a simple distributed mobility scheme (*Random Movement* or RM) according to which the sink moves uncontrolled and randomly throughout the network. The different mobility schemes are compared through extensive ns2-based simulations in networks with different nodes deployment, data routing protocols, and constraints on the sink movements. In all considered scenarios, we observe that moving the sink always increases network lifetime. In particular, our experiments show that controlling the mobility of the sink leads to remarkable improvements, which are as high as sixfold compared to having the sink statically (and optimally) placed, and as high as twofold compared to uncontrolled mobility.

Keywords Wireless sensor networks · Controlled mobility · Mobile sensor networks

S. Basagni (✉)
Department of Electrical and Computer Engineering,
Northeastern University
e-mail: basagni@ece.neu.edu

A. Carosi · C. Petrioli
Dipartimento di Informatica, Università di Roma "La Sapienza"
e-mail: carosi@di.uniroma1.it

C. Petrioli
e-mail: petrioli@di.uniroma1.it

E. Melachrinoudis · Z. M. Wang
Department of Mechanical and Industrial Engineering,
Northeastern University
e-mail: emelas@coe.neu.edu

Z. M. Wang
e-mail: zmwang@coe.neu.edu

1 Introduction

Recent years have witnessed an increasing interest in wireless sensor networks (WSNs). These networks are made up of wireless nodes endowed with sensing capabilities that are deployed for implementing a host of different applications. Typical examples of WSN applications include environmental monitoring, independent assisted living, disaster assessment and recovery, control of industrial processes, etc. [3, 8].

From a networking perspective, WSNs generally follow the well-established *ad hoc* paradigm of communication: Data delivery between any two nodes follows a multi-hop route. Differently from ad hoc networks, where any two nodes can be source and destination of data packets, in WSNs data generated by the sensor nodes are sent to one or more data collection points (the *sinks*). Sinks are considered resource-rich, i.e., energy, processing power and memory are not considered a limitation for their prolonged functioning and operations. Sensor nodes are instead usually quite constrained in terms of battery power, storage, and computational capabilities. Due to the large number of sensor nodes deployed and to their being often placed in hostile, unaccessible environments, it is not viable to recharge/replace their batteries. Whenever a node depletes its energy it is considered "dead," i.e., no longer able to perform its sensing or communication duties.

Several protocols have been proposed so far for data delivery and dissemination in WSNs. One of the main aims of these solutions is minimizing the nodes energy consumption (mostly due to radio communications) in order to increase the time that the network is able to perform its intended operations (*network lifetime*). Independently of all the energy-efficient techniques developed at the different layers of the nodes protocol stack the ultimate problem concerns the delivery of the sensed data from all the sensors to the sink, which imposes greater burden on nodes that are closer to the sink. More specifically, when a sink is statically placed, the sensor nodes that can directly communicate with it (the sink's *neighbors*) tend to deplete their energy faster than other nodes. Not only they consume energy to communicate their own data to the sink, but also for relaying to it the data from any other node. This problem, here termed the "sink neighborhood problem," leads to a premature disconnection of the network. The sink gets isolated from the rest of the network due to the death of its neighbors while most of the sensor nodes are still fully operational.

One way for mitigating—if not obviating—the sink neighborhood problem is by exploiting the mobility of some of the network components. The key idea is that of changing the neighbors of the sink so that the energy consumption for data packet relaying is balanced throughout the network. Since moving the nodes would require extra power from the already limited energy of a node, the most promising way of changing the sink's neighbors is to have the sink moving to different parts of the deployment area, while keeping the sensors static.

Protocols proposed so far for sink mobility differ on the nature of the mobility itself. *Uncontrolled* sink mobility is used in those applications where the sink is sent to gather data through the network at times and along routes that are out of the control of the network. Whether random or deterministic, the sink movement proceeds according to a schedule which is not determined by the prevailing network conditions, such as data traffic or the nodes residual energy.

More recently, several protocols have been proposed that show how, by having the network *controlling* the mobility of the sink, remarkable improvements can be obtained, especially for extending network lifetime [14, 31, 35, 53]. All these solutions are centralized, in the sense that the proposed schemes determine optimal sink routes and sojourn times based on the knowledge of global network parameters.

This paper contributes in multiple ways to the investigation of using controlled sink mobility for extending WSN lifetime.

We start by presenting a new Mixed Integer Linear Programming (MILP) model that determines sink routes and sojourn times at the *sink sites* (specific locations the sink can visit). Differently from previous solutions, we include parameters and constraints that model realistic requirements of a WSN. For instance, we consider the cost of moving the sink from a site to another, both from a data latency point of view (as data packets need to be buffered during the sink movement) and from an energy consumption point of view (we consider the cost of building and releasing data routes from the sensors to the current position of the sink explicitly). We also introduce constraints for considering the mobility rate of the sink, imposing a minimum sojourn time for the sink at the different sites. This allows us to investigate how lower or higher sink mobility affects network lifetime. Although data routing optimization could be easily incorporated in our MILP model, we choose it to be routing-independent. First of all, a centralized solution for routing is not viable for WSNs. Moreover, deriving routing as part of the model solution optimizes it only with respect to the sole metric of network lifetime, whereas there are other relevant metrics to be considered when designing efficient routing for WSNs (e.g., packet throughput, data latency, control overhead, etc.).

ILP-based solutions are notoriously hard to compute [15], and often times these models can be solved only for particularly small input scenarios. However, the simplicity of our new model for sink mobility makes it possible to obtain optimal sink routes for non-trivial, quite realistic cases, such as networks with hundreds of sensor nodes and several dozens of sink sites. This allows us to use the MILP-generated optimum sink mobility as an upper bound for more suitable distributed and localized solutions to be devised and deployed.

MILP models provide *centralized solutions*. For instance, in order to find lifetime-optimal routes and sojourn times for the mobile sink one has to provide a global view of network topology, communication costs, etc. Centralized solutions, however, are unbearably time and energy consuming for most WSNs applications.

The second contribution of this paper concerns the development of decentralized, simple solutions. To this aim, we introduce a completely distributed and localized protocol for sink mobility. Throughout the network lifetime the sink moves as drawn by sink sites in energy-rich areas of the network. More specifically, the sink keeps monitoring surrounding sites with respect to the energy of the nodes around them. When a site different from the current is in an area with higher energy, the sink greedily moves at that new site. This simple heuristic, termed *Greedy Maximum Residual Energy* (GMRE), takes explicitly into account crucial parameters such as the costs of data route release and establishment when the sink moves, different sink mobility rates, as well as possible constraints on sink mobility. It is also completely adaptive to different data routing protocols. To the best of our knowledge, this is the first completely distributed and localized solution for sink mobility in WSNs.

The third contribution of this work concerns a thorough comparative performance evaluation of mechanisms for sink mobility. To this aim, we have performed an extensive set of ns2-based simulation experiments that compare the performance of GMRE to that of the following protocols.

(1) The omniscient optimal solution provided by our MILP model. Once crucial parameters are properly tuned we observe that GMRE achieves average network lifetimes that are only 22% shorter than the optimal ones in the worst case. The average decrease in network lifetime of GMRE vs. OPT is only 13.4%. Given that according to GMRE the sink makes movement decisions based on the local knowledge of the network, we consider this a promising, remarkable result.
(2) A simple heuristic, here termed Random Mobility (RM), that represents the uncontrolled sink mobility approaches often seen in the literature. In all considered scenarios GMRE obtains average network lifetimes that are from 50 to 100% longer than those achieved by RM.
(3) The case when the sink is statically and optimally placed at the center of the deployment area. While the optimum achieves improvements up to 500% with respect to the static placement, and RM is able to provide longer lifetimes (about 200% longer lifetimes), we show that GMRE is always closer to the optimum, being able to prolong network lifetime up to 4 to 5 times.

Our simulations go beyond a quantitative assessment of the performance of the compared mobility schemes. Our aim is to gain a clear understanding of the rationale beyond the obtained results and a deeper comprehension of the impact on network performance of protocol parameters and different scenarios.

The experiments are organized in three parts. At first we select a basic testing scenario to provide clear evidence of the advantages of sink mobility and the effectiveness of a greedy approach to it for longer network lifetime. We then proceed by showing the effect of key protocol parameters over the tested protocols. More involved scenarios, which comprise different node deployment schemes, different data routing protocols and sink site locations are finally considered in our third set of experiments. Our experimental investigation aims to more than showing betterment in network lifetime. Other important metrics of interest, such as data latency, protocol overhead and energy consumption over time are also evaluated. For all the metrics considered, we provide both quantitative results as well as an in-depth explanation of the motivations behind the protocols behavior. The GMRE heuristic is shown to provide good trade-offs in terms of all the considered metrics.

The paper is organized as follows. In the next section we review previously proposed solutions that exploit the mobility of different network components (e.g., sink, relay nodes and sensor nodes) according to different degrees of mobility control. Section 3 defines in detail the problems of controlled sink mobility and the scenarios we consider. The following Section 4 defines the MILP formulation, outlining also its novelty, generality and strengths. The new distributed heuristics introduced in this paper are described in Section 5. Simulation results are presented and explained in Section 6, where we also draw conclusions on the comparison between the different sink mobility schemes. Section 7 concludes the paper.

2 Related works

Considering mobility as a "blessing" rather than a curse for network performance has been widely discussed for ad hoc and sensor networks in different contexts [11, 12, 17, 23, 26, 27, 30, 34, 56, 57]. The primary objective of these works is to deliver messages in disconnected ad hoc networks and to improve network throughput.

The work by Chatzigiannakis et al. [11] explores the possibility of using the coordinated motion of a small number of users in the network to achieve efficient communication between any pair of other mobile nodes. A fraction of the network nodes acts as forwarding agents carrying packets for other nodes: The packet is exchanged when the source node and the agent are neighbors (i.e., in the radio vicinity of each other), and it is then delivered to the intended destination when the agent passes by it.

This basic idea has been introduced to WSNs by Shah et al. in their works on *data MULEs* [23, 38]. Mobile nodes in the sensor field, called MULEs, are used as forwarding agents. The idea here is to save energy by having single-hop communication (from a sensor to the MULE that is passing by) instead of the more expensive multi-hop routing (from the sensor to the sink): It is the MULE that will eventually take the sensed data to the sink. The data MULE architecture is effective for energy conservation in delay tolerant networks [39]. Energy is traded off for latency, i.e., the energy needed to communicate a packet to the sink is decreased at the cost of waiting for a MULE to pass nearby (and at the cost of waiting for the mule to move to the vicinity of a sink). Scheduling problems in sensor-to-sink transmissions within this model have been studied in [41].

A dissemination protocol where a tree-like communication structure is built and maintained is described by Kim et al. in [29]. The randomly moving sinks access the tree from specified sensor nodes in the tree (access points). Communication between the sink and the access points can be multi-hop. This happens when the sink moves away from the access points. The protocol, termed SEAD (Scalable Energy-Efficient Asynchronous Dissemination), is designed for finding a trade off between the end-to-end packet delay and the energy consumption spent for reconfiguring the tree so that the access points are closer to the current position of the sinks. SEAD is shown to be more effective for conserving energy than other solutions for data dissemination in wireless sensor networks such as directed diffusion [22], TTDD [55] and ADMR [25].

The problem of building and maintaining routes to a mobile sink with the precise aim of minimizing the corresponding overhead has been tackled with in [20, 21] and [2]. In the first paper, local update techniques are described for detecting disconnections and for performing route repair in "sink oriented trees." In [20], the ERUP protocol is proposed for conducting route re-discovery only in the vicinity of the damaged route. In [2], initial routes are constructed from the nodes to the sink according to any viable WSN routing. If, because of the movement of the sink, the routes are no longer valid, forwarder nodes are designated to extend the current routes.

Common to all these works is that the mobility of the sink is unpredictable and *uncontrollable*. For example, in [29] and [54] sinks move according to the random waypoint model.

The use of mobile sinks with *predictable* mobility has been more recently presented in [32, 45, 46] and [9]. In these works the sinks (airplanes) fly over the sensor field and gather the sensed data periodically. While the movement of the sink is fully controllable, it is external to the network infrastructure, i.e., the trajectories are not determined by network components and activity. The main contribution of these papers concerns the energy-efficient transmission to the passing sink [45–47]. In [32] the authors consider heterogeneous sensor networks made up of two types of nodes, and determine the densities of each type and the battery energy needed to achieve a given network lifetime.

Inherent patterns of the sink movement are exploited in [4] for the design of robust and energy-efficient routing. This paper assumes that there is a certain degree of predictability in the sink movement, such as the routine route of a ranger patrolling a forest. Based on statistics and distributed reinforcement learning techniques, the sensor nodes learn about the sink whereabouts at given times and use this information to find routes to the mobile sink.

A model for sink movement is proposed in [9], where "observers" (i.e., the sinks) move along the same path repeatedly. The sensed data are collected while the observer traverses the network. When passing by sensor nodes, the observer wakes them up and receives their data (if any). The authors describe a prototype system developed at Rice University where the observers are carried by campus shuttles, and the sensors are spread out throughout university property. In particular, the authors determine the transmission range needed to collect data from a predefined percentage of the sensor nodes, given the observer speed, the time required to transmit a piece of information, and the traffic pattern. The correlation among the various system parameters is investigated analytically.

The idea of using unmanned vehicles as data collectors has been further investigated in [44]. The sensor nodes send their data to nearby clusterheads via multi-hop routing. The vehicles then pass by the clusterheads to collect the data. Three different solutions are presented in the paper that define different schedules for the collectors to visit the clusterheads. In this way, the collector has to visit only the clusterheads, and not all the nodes, while multi-hop routing is reduced to a smaller number of hops since the data is sent from a sensor node to a clusterhead who is nearby. In this case the sink (base station) is not really mobile, and the collectors return to the base station periodically to deliver the data and for recharging. The architecture of [44] has been recently expanded in [43] to consider different classes of nodes, where the collectors roam (controllably) among the clustered sensor nodes some of which can be (uncontrollably) mobile. Aim of the paper is to determine schedules for the collectors to visit the nodes so that transmission energy, data latency and buffer requirements at the nodes are minimized.

In this paper we are interested in techniques for prolonging the lifetime of a wireless sensor network, and in particular in how *network-controlled* mobility can improve data dissemination and collection, especially from an energy consumption perspective. There are three major approaches to the use of controlled mobility in WSNs: The sink itself moves among the sensor nodes and collects the data; mobile relays are used for data gathering and following delivery to the (static) sink; and finally, the sensor nodes are mobile.

The first two approaches appear to be the more promising for energy efficiency and longer network lifetime, since sink and relays are usually considered resource-rich, and hence energy consumption and network lifetime are not impacted by the energy needed to move them. In the case the sensor nodes move, a great deal of the nodes' energy is spent on the movement itself, thus having a detrimental impact on the nodes' lifetime. Works that consider mobile sensors are mostly concerned with sensor deployment time and sensing coverage. The costs associated with sensor movements as well as the cost of transmitting sensed data are often not considered, and network lifetime is rarely a metric of interest. For results in this area, which goes beyond the scope of this paper, the reader is referred to [16, 18, 19, 36, 48–51].

A first in-depth discussion on how to incorporate controllable mobile *relays* into the network infrastructure has been presented in [28]. The authors describe an implementation of a sensor network with an autonomous mobile relay (a robot) that visits the (static) sensors, collects their data, and delivers them to the sink. The idea of collecting data in a single-hop fashion (i.e., when the robot approaches a sensor) is similar to that of data mules. However, in this case the movements of the robot adapt to data collection performance parameters, which are dictated by the network application priorities. The robot is part of the system, and it is the system that controls its mobility. The testbed-based experimental results in this paper concern the evaluation of methods for controlling the speed of the robot for optimizing data collection. The robot traverses networks with different densities following a straight trail and collects the data that are then brought to the sink. The authors further explore the use of controlled relays in [24, 40]. The problem of controlled data MULE-like mobility has been also recently addressed in [13]. The authors first propose an algorithm for avoiding sensor nodes buffer overflow while minimizing the speed of the mobile relay. They also extend this algorithm to the case where some of the packets have delay constraints (i.e., they are "urgent messages" that have to be delivered to the sink within a given time since their generation).

An investigation of the controlled use of relay nodes for data collection and subsequent report to the sink has been recently proposed in [52]. Although the authors recognize that moving the sink directly yields better resource utilization and hence longer lifetime, they argue that for certain applications moving the sink might be infeasible. Therefore, having one or more resource-rich mobile relay nodes is remarkably helpful. Given that the sensor nodes know about the current location of the relay node, routing protocols are presented for delivering the data from the sensors to the relay, from the relay to the sink, and finally for determining the route of the relay. Improvements on network lifetimes are fourfold with respect to the case of a static sink.

Works on performance of mobile relays show that moving the sink appears to be more promising for achieving better trade offs between energy consumption and data latency, motivating research on mobile sinks. In this field, results on reducing energy consumption and on the maximization of the lifetime of a sensor network has been tackled with in [14] and, more recently, in [35, 53] and [31]. In these works, centralized algorithms are presented where the sink moves among the (static) sensor nodes and, while sojourning at given locations, collects data that are sent to it via multi-hop routes. The first work is mostly concerned with energy minimization. The authors present an ILP model to determine the locations of multiple sinks and the routes from the sensors to the sinks. Time is divided into rounds. At the beginning of each round information on the nodes' residual energy is centrally gathered and the ILP problem is solved to determine new, feasible locations the sinks should travel to for minimizing the maximum energy consumption spent at the nodes during that round. Minimizing the energy consumption results in increased network longevity. No constraints are enforced on the sink movements, and there is no relation between the number of the sinks and their position in subsequent rounds.

We have explicitly addressed the problem of network lifetime maximization through controlled sink mobility in [53] for networks with a single sink. Sink locations (in this case the sensor sites) and sink sojourn times at those locations are determined that maximize the network lifetime via a new LP formulation of the problem: Maximizing the network lifetime equals maximizing the sum of sojourn times of the sink at the visited locations. Although the model is completely general, the experiments performed in the paper refer to scenarios where $n = L^2$ nodes are arranged in a $L \times L$ grid. The sink has no limitation on the time $t_k \geq 0$ it can spend at sensor k and can move from any location to any location in the network. Improvements on network longevity are obtained that are almost five-fold when the sink sojourns at the nodes located at the four corner areas and in the central area of the grid.

By combining the model presented in [53] and the LP formulation for maximum lifetime routing described in [10], Papadimitriou and Georgiadis [35] present another centralized solution for the problem of maximizing network lifetime. By turning a constant of the model in [53] into a variable, the model presented in [35] jointly solves the problem of determining the sink sojourn times at the given sites, and the routing of the packets to the current position of the sink. This (data) routing-dependent solution achieves improvements with respect to the lifetime values of [53] that are twofold.

The problem of lifetime maximization has been formulated as a min-max problem by Luo and Hubaux [31]. By considering together sink mobility and data routing, a load balancing solution is obtained that, while keeping the sink moving along the external perimeter of the network, achieves lifetimes 500% higher than when the sink stays in the center of the network. Among the contributions of this work it has to be noticed an explicit reference to the complexity of formalizing the optimization problem of maximizing network lifetime via sink mobility with traditional analytical tools.

Finally, the problem of devising distributed solutions for network controlled sink movements has been recently tackled with by the authors in [5, 6], where preliminary, promising performance results on network lifetime are shown for the solutions presented in this paper.

3 Problem definition

We consider a scenario where a large number $|N|$ of resource constrained, static nodes with sensing and wireless communication capabilities are scattered in a given geographic area. (In the following these nodes are referred to as *sensor nodes*,

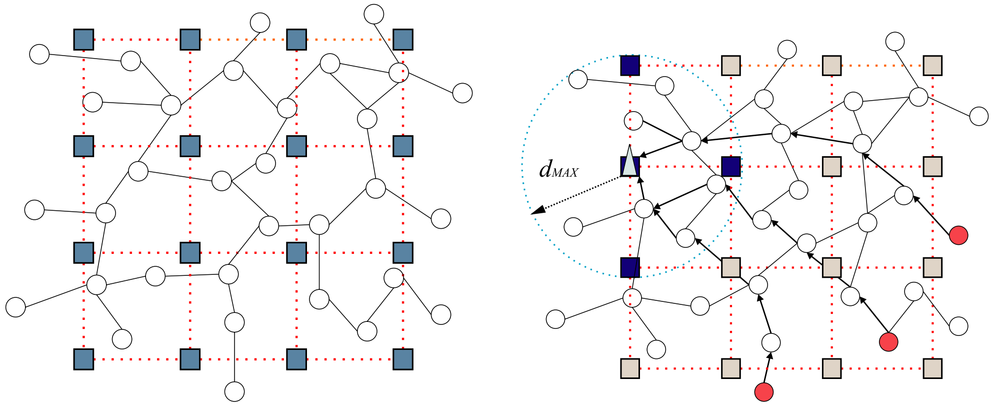

(a) Sensor nodes and sink's sites (b) Sink movements and routes to the sink

Fig. 1 Typical WSN scenarios

or often, simply, nodes.) We assume periodic generation of data packets at the sensor nodes: Node $i \in N$ transmits at a given data rate r_i packets that are "convergecasted" to the sink for processing. While the nodes are static, the sink can be mobile. More specifically, we consider a set $S = \{1, \ldots, q\}$ of q sink's *sites* which are the points within the geographic area the sink can visit. For instance, Fig. 1(a) shows a typical scenario where 32 nodes (represented by circles) are scattered randomly and uniformly on a bi-dimensional area, and 16 sink's sites (squares) are organized according to a 4×4 grid. A (solid) link between two nodes indicates that those two nodes are neighbors (i.e., they can hear each other's transmissions). A (dotted) link between two sites indicates that the sink can move from one site to the other and vice-versa.

Because of the sink's neighborhood problem the sink moves throughout the network in an attempt to balance the energy consumption among the nodes. Every time the sink reaches a new site, it floods a packet f to all the network nodes making them aware of its current site. A node that receives f starts sending/relaying its packets toward the new site of the sink. Every routing scheme that works with the topological information provided by f, such as geographic or shortest paths-based routing, is a viable routing for data delivery to the sink. We observe that the independence from the particular routing protocol yields a twofold advantage. First of all, it guarantees the longest possible network lifetime given the specific routing. Furthermore, it allows the network users to design or choose the routing algorithm that best meets the WSN applications requirements in terms of a host of different metrics of interest (not just the lifetime). Every time the sink leaves a site, it again floods a packet to all nodes to communicate that it is no longer reachable at that site. Upon receiving this second packet, a node stops forwarding data (remaining packets are buffered), and waits to receive a new packet f from the sink, carrying its whereabouts. When the new packet f arrives and routes to the new site of the sink are formed, packet transmission is resumed.

There is virtually no bound on how far the sink can travel between two sites. However, we note that while the sink is traveling, the sensors do not transmit. Therefore, if new data are sensed, these are buffered. This implies the possibility of high delays for data packets. In order to contain this delay, we introduce a new parameter d_{MAX} which represents an upper bound on the distance that the sink can travel from a site to the following one. Thus, the pair (S, d_{MAX}) uniquely defines a graph of sink's sites where there is a link between two sites if and only if their (Euclidean) distance is $\leq d_{\text{MAX}}$. Figure 1(b) shows the three sites (darker squares) the sink (the triangle) can reach from its current position. The dotted lines between the sites of Fig. 1 indicate that the sink can only move horizontally or vertically in the 4×4 grid. (Routes from selected sensors to the current site of the sink are also shown.)

We observe that in case of high sink mobility and low data traffic the energy cost for route construction and release can be significant. Therefore, this cost is explicitly taken into account. In order to evaluate the impact of different (higher or lower) sink mobility rates, we introduce the parameter $t_{\min}$ to represent a mandatory minimum time the sink has to sojourn at a site.

We will solve the following problem via mathematical modeling and by designing distributed protocols:

> *Determine the starting site and the route for the mobile sink over the graph* (S, d_{MAX})*, together with the sojourn times* $t_k \geq t_{\min}$ *of the sink at each visited site* $k \in S$ *so that network lifetime is maximized.*

4 Mathematical model

In this section we present a Mixed Integer Linear Programming (MILP) formulation of the problem described above. We start by defining the sets, the parameters and the variables used for formalizing our problem.

Sets and parameters

- S is the set of sink's sites, i.e., the locations at which the sink may sojourn: $S = \{1, \ldots, q\}$.
- N is the set of the network nodes: $N = \{1, \ldots, n\}$.
- e_0: Initial energy (Joules) of each node.
- f_{ik}: Energy consumption (Joules) at node $i \in N$ for setting up/releasing routes when the sink moves to site $k \in S$.
- c_{ik}: Power consumption (Watts) for receiving and transmitting packets at node $i \in N$ when the sink sojourns at site $k \in S$.
- $t_{\min}$: Mandatory minimum time (secs) for which the sink is required to stay at site $k \in S$.
- d_{jk}: Euclidean distance (meters) between any two sink sites $j, k \in S$.
- d_{MAX}: Maximum distance (meters) the sink is allowed to travel each time it moves.
- A: The set of directed edges joining sink sites whose distance is less than or equal to d_{MAX}, i.e., $A = \{(j,k) \in S \times S : j \neq k, d_{jk} \leq d_{\text{MAX}}\}$.
- O: The set of directed edges $(0, k)$, $k \in S$, joining a fictitious site 0 (origin) with the sites in S.
- D: The set of directed edges $(k, q+1)$, $k \in S$, joining the sites in S with a fictitious site $q+1$ (final destination).
- X: The union of A, O and D.

Variables

- t_k: Sojourning time (secs) of the sink at site $k \in S$.
- y_k: Binary variable taking the value 1 if the sink sojourns at site $k \in S$ ($t_k > 0$); 0 otherwise ($t_k = 0$).
- x_{jk}: Binary variable indicating the status of $(j,k) \in X$. $x_{jk} = 1$ if and only if arc (j,k) is on the sink movement route; $x_{jk} = 0$ otherwise.
- u_k: Auxiliary variable used to enforce a unique sink path.

MILP formulation

$$\text{Max} \sum_{k \in S} t_k \tag{1}$$

$$\text{subject to:} \quad \sum_{k \in S} c_{ik} t_k + \sum_{k \in S} f_{ik} y_k \leq e_0 \quad (i \in N) \tag{2}$$

$$t_{\min} y_k \leq t_k \leq M y_k \quad (k \in S) \tag{3}$$

$$\sum_{k \in S} x_{0k} = 1 \tag{4}$$

$$\sum_{k \in S} x_{k,q+1} = 1 \tag{5}$$

$$\sum_{\substack{j \in S \cup \{0\} \\ (j,k) \in O \cup A}} x_{jk} = \sum_{\substack{j \in S \cup \{q+1\} \\ (k,j) \in A \cup D}} x_{kj} \quad (k \in S) \tag{6}$$

$$\sum_{\substack{j \in S \cup \{0\} \\ (j,k) \in O \cup A}} x_{jk} = y_k \quad (k \in S) \tag{7}$$

$$u_j - u_k + q x_{jk} \leq q - 1 \quad ((j,k) \in A) \tag{8}$$

$$t_k, u_k \geq 0 \quad (k \in S) \tag{9}$$

$$y_k \in \{0, 1\} \quad (k \in S) \tag{10}$$

$$x_{jk} \in \{0, 1\} \quad ((j,k) \in X) \tag{11}$$

The objective function (1) maximizes the sink's total time at sojourning sites, $\sum_k t_k$, which is the effective network lifetime. Constraint (2) states that the combined energy spent at node i for data delivery ($\sum_{k \in S} c_{ik} t_k$) and for data route construction and release ($\sum_{k \in S} f_{ik} y_k$) during $\sum_k t_k$ (the time before the death of the first node) should not exceed the node's initial energy e_0. The right part of double inequality (3) forces y_k to take the value 1 if the sink sojourns at site k ($t_k > 0$), thus linking the binary variable y_k (constraint (10)) with the continuous variable t_k. M is a significantly large number. The left part of double inequality (3) restricts the sojourn time t_k to be at least equal to the mandatory minimum sojourn time $t_{\min}$ if the sink sojourns at site k ($y_k = 1$) and at the same time forces y_k to take the value 0 if the sink does not sojourn at site k ($t_k = 0$). The first sojourning site in the sink's movement route is allowed to be any site in S. To implement this, a fictitious fixed initial site 0 (origin) is introduced. At the beginning of the sensor network's lifetime, the sink moves in zero time (and cost) from the origin to some site $\alpha \in S$, determined by the model. This is that particular site such that $x_{0\alpha} = 1$ (Eq. (4)), namely, it is the optimum starting point of the sink journey. Then, the sink sojourns at that first site and at subsequent other sites in S to be determined by the model. Finally, from the last sojourning site ω the sink moves to a second fictitious site "$q+1$" (destination), again in zero time (Eq. (5)). The site ω completes the sink route started at site α. This is the last site at which the sink sojourns, and marks the end of the sensor network lifetime. The arcs $(j,k) \in X$ on the sink route are associated with binary variables x_{jk} equal to 1. The variable x_{jk} is equal to 0 for all the $(j,k) \in X$ that do no belong to the route. Equivalently, one can think of a unit of flow moving from the origin to the destination. Constraint (4) induces a unit of flow from the origin to some node $\alpha \in S$, while constraint (5) causes the destination to absorb a unit of flow coming from some node $k \in S$. Constraint (6) forces flow conservation at all sites $k \in S$, thus ensuring the generation of a route. Constraint (7) ensures that the sites $k \in S$ on the generated route are sites at which the sink sojourns ($k | y_k = 1$). To elaborate, if y_k in constraint (7) equals 1, then the sink sojourns at site k, and therefore there must be one and only one arc on the sink's movement route reaching site k. On the other hand, if y_k equals 0, then there will not be any incoming arc to that site.

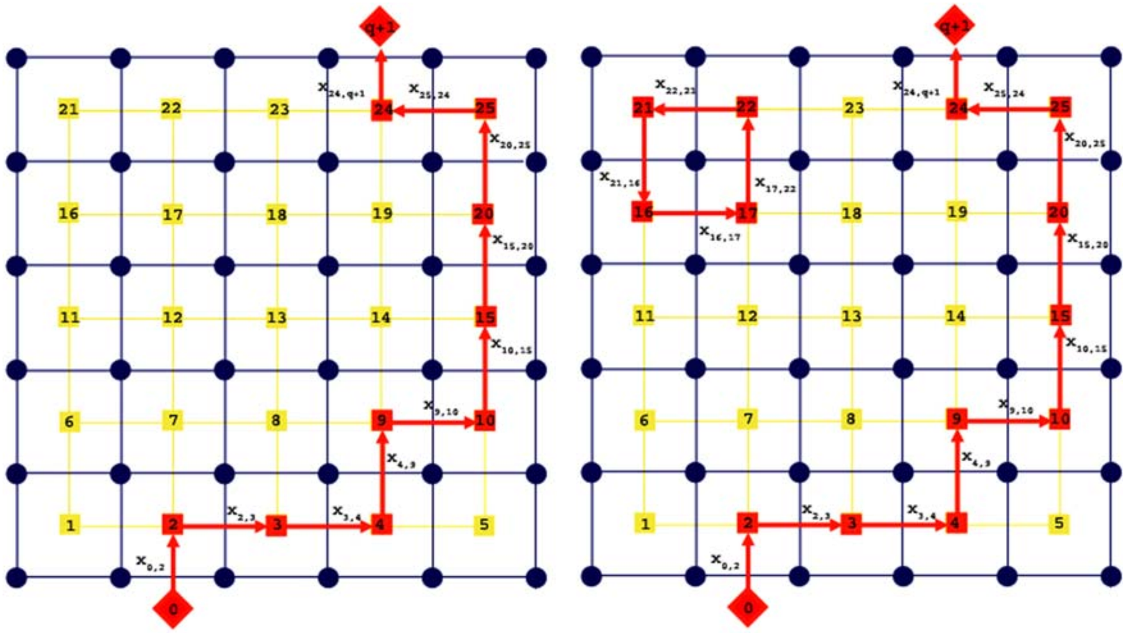

(a) A sink route (b) A disjoint cycle in the sink route

Fig. 2 Sink optimum routes produced by constraints (4) to (7)

Finally, constraint (5) induces a unit of flow from the last node in the sink route (ω) to the fictitious final node $q+1$.

Figure 2(a) shows a possible optimum sink route that goes from the initial site $\alpha = 2$ to the final site $\omega = 24$. Constraints (4) and (5) ensure that, independently of α and ω, respectively, there is only one initial site and one final site for the sink route, since the corresponding arcs ($x_{0\alpha}$ and $x_{\omega q+1}$) must be 1. The combination of constraints (6) and (7) takes care of generating the route between α and ω that passes through all the sites where the sink has to sojourn for maximizing the network lifetime. In particular, the first constraint mandates that there must be one outgoing arc x_{kl} for every incoming one x_{jk} (with the natural exception of the two fictitious sites 0 and $q+1$). For instance, this is the case of arcs x_{34} and x_{49} in Fig. 2(a), which are both set to 1. The remaining arc out of site 4, namely, arc x_{45}, is forced to be 0. According to constraint (7) for every site k where the sink sojourns ($y_k = 1$) there must be a way to get there, i.e., there must be exactly one site j (which includes the fictitious site 0) from which k is reachable ($x_{jk} = 1$). At the same time, the sink should not pass through sites where it does not sojourn. For instance, sites with $k = 2, 3, 4, 9, 10, 15, 20, 25$ and 24 in Fig. 2(a) are all and only those for which $y_k = 1$, i.e., these are all and only those sites that can be in the sink route. All other sites h are such that $y_h = 0$.

We note that flow conservation constraints (6) and (7) do not prevent the formation of cycles disjoint from the route from the origin to the destination. The (disjoint, non simple) route depicted in Fig. 2(b) comprising nodes 2, 3, 4, 9, 10, 15, 20, 25 and 24 and nodes 16, 17, 22 and 21 (cycle) is possible according to our model up to constraint (7), since none of these constraints is violated by having $y_{16} = y_{17} = y_{21} = y_{22} = 1$ as well as $x_{16,17} = x_{17,22} = x_{22,21} = x_{21,16} = 1$. This situation is undesirable, since quite unrealistic. It is practically impossible, for instance, to have the sink moving from site 9 (a site in the connected route from site 2 to site 24) to site 17 (a site in the cycle): Sites 9 and 17 are not directly connected, i.e., their distance is $\geq d_{\mathrm{MAX}}$. Constraint (8) ensures that no such cycles are formed. (A similar constraint has been used in the integer programming formulation of the Traveling Salesman Problem (TSP) to avoid sub-tours [33].) According to constraint (9) a site k is associated with a "weight" $u_k \geq 0$. Constraint (8) imposes that the sites visited by the sink are traversed in increasing order, i.e., if $x_{jk} = 1$ then $u_j < u_k$. This renders clearly impossible to return to the same node, and hence to form cycles like the one in Fig. 2(b).

Some comments are in order. The parameter $t_{\min}$ has been introduced to assess the effect of different (higher or lower) sink mobility rates on network performance. For a given $t_{\min}$ the model will produce the sink route and sojourn time $t_k \geq t_{\min}$ at site k that maximizes network lifetime. By varying $t_{\min}$ we can explore a number of trade-offs. For instance, at higher $t_{\min}$s we expect to have lower overhead (e.g., for route construction and release). Shorter $t_{\min}$s result in better choices of sojourn times at different sites (which might result in a longer network lifetime) at the price of increasing overhead. Depending on prevailing network conditions, there is a value of $t_{\min}$ for which the advantage of a finer granularity of sojourn times is outpowered by the energy consumption increase due to the extra overhead.

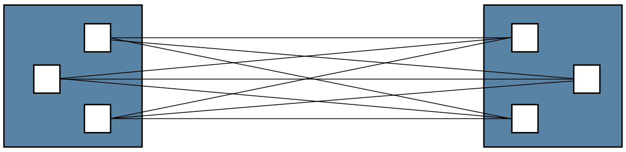

Fig. 3 Two adjacent physical sites, six logical sites and their interconnections

The model is flexible in letting the user employ any method for determining the power consumption rate c_{ik} of each node $i \in N$ when the sink sojourns at site $k \in S$. This cost depends on both node i's transmission rate r_i and on the particular protocol for routing the packet to the sink sojourning at site k. The costs c_{ik}s could be computed analytically [53], or they can be provided as input to the model from simulations [5] or from real-data traffic traces. In short, the model can be customized to find the optimum lifetime for different routing protocols (by computing the corresponding values of f_{ik} and c_{ik}).

Constraint (8) renders infeasible all cycles formed by the nodes in S, thus allowing only a unique simple path. We notice, however, that the model can be easily extended to allow the sink to sojourn at the same site multiple times. A single "physical" site can be represented by h "logical" sites, where h is the number of times we want the sink to be able to pass through that site. The logical sites have no arcs between them, and are connected to all the (logical) sites of adjacent (physical) sites. Figure 3 depicts the case of two adjacent physical sites and the corresponding six logical sites ($h = 3$), along with their interconnections.

With this simple modification we obtain optimal lifetime given that the sink is allowed to visit each site at most h times. The global optimum is obtained by progressively running the solver of the model on increasing hs until the lifetime stabilizes.

Our MILP formulation improves over previously proposed models in multiple ways. The model is independent of a number of factors such as the specific sensor node deployment and sensor density; the sink site topology; the size and shape of the geographic area of deployment, and the sensor node technical features (e.g., transmission radius, energy model, etc.). The given formulation includes a number of realistic constraints, such as the non-instantaneous movement of sink between sites potentially far apart from each other. Most importantly, and differently from all previously proposed LP solutions, our formulation explicitly includes the costs for changing sink sites.

5 Distributed protocols

In this section we describe the details of the two new distributed protocols for sink mobility that we are going to compare with the optimal routing strategy provided by the MILP model.

The two protocols differ on the strategy used by the sink sojourning at a site for choosing the next one.

In the Greedy Maximum Residual Energy (GMRE) protocol the sink greedily selects the site within d_{MAX} surrounded by nodes that have the most energy left. The idea is that in time, this should most likely result into a balanced energy consumption throughout the network, and hence into a longer network lifetime. After spending a time $t_{\min}$ at a site, a sink evaluates whether to move toward one of the adjacent sites or to stay where it is. Two sites are adjacent if their distance is $\leq d_{\text{MAX}}$. In order to decide whether to move or not, the sink gathers information about the residual energy at the nodes around each of the potential future sites (we call this energy value the residual energy at the site), and compare it with the residual energy at the current site. If there are adjacent sites with a residual energy higher than that at the current site, the sink moves to the site with the highest residual energy (selecting randomly among sites with the same residual energies in case of ties). Otherwise the sink stays at the current location.

Key to the definition and implementation of GMRE is the communication to the sink of the residual energies at the adjacent sites. This communication proceeds in two phases. First, for each of the adjacent sites, the sink identifies one *sentinel* sensor node that will be in charge of measuring and reporting the residual energy at that site when requested by the sink. The second phase concerns the sink interrogation of the sentinels. This is performed whenever the sink has to decide whether to move or not.

To implement the first phase we take advantage of the flooding performed by the sink when it makes the nodes aware of its new location. For this heuristic protocol we assume that a node that is in the "transmission vicinity" of a site (i.e., whose Euclidean distance from a site is less than or equal to the nodes transmission range) is aware of that. This can be obtained by endowing the nodes with a suitable localization mechanism (such as one of those described in [37]). The flooding message contains the coordinates of the current location of the sink. Upon receiving the flooding packet, a node knows if it is in the vicinity of a possible future sink site. In this case, it sends to the sink a (small) packet for its candidacy as sentinel. Upon receiving such packets the sink decides which is the sentinel for a given site. This mechanism also allows the sink to identify those sites that are isolated (no packet is received from nodes around that site). In this case, the sink will not consider that site as a possible future one.

The second phase starts when the sink has to decide whether to move to a new site or not. At this time, the sink interrogates the selected sentinels about the residual energy at their sites. This is accomplished by sending a (small) packet to the sentinels. When interrogated, the sentinels query their neighboring sensor nodes about their residual energy and communicate back to the sink the minimum of the obtained values (or any suitable function that can express how critical for the network lifetime is to place the sink in that area).

The second, simple protocol for sink mobility that we propose here captures the case of uncontrolled, random mobility of the sink. Every $t_{\min}$ the sink selects randomly and uniformly the new location among all the sites within distance d_{MAX} from the current. In case a site different from the current is selected, the sink moves to that site. This simple scheme, referred to in the following as the Random Movement heuristic (*RM*), generalizes data gathering protocols previously proposed in the literature (e.g., the data mules approach [23]) to the case of multi-hop data routing. We use RM mainly as a benchmark for assessing the effectiveness of GMRE in prolonging network lifetime.

6 Simulation results

In this section we discuss the results of a thorough ns2-based [42] performance evaluation of the presented protocols. In particular, we have compared the following four sink mobility schemes.

(1) The sink is static. This is a degenerate mobility scheme. So degenerate, in fact, that the sink does not move. In this case, that we name STATIC, the sink is placed at the geographical center of the deployment area, which is the position that maximizes the network lifetime.
(2) The sink moves along the optimum route derived by the MILP model presented in Section 4 (OPT mobility in the following).
(3) The sink moves according to the RM scheme.
(4) The sink moves as specified by the GMRE heuristic described in Section 5.

The performance of the four schemes have been compared with respect to the following metrics.

- *Network lifetime*, i.e., the time until the first node dies having fully depleted its energy.
- *Per node residual energy over time*. Investigating this metric allows us to determine the actual energy consumption associated with the different mobility schemes as well as to verify how balanced is the consumption itself.
- *End-to-end packet latency*. This is the time that goes from packet generation at a sensor node to the successful delivery of that packet at the sink.
- *Overhead* (bits/s). The overhead incurred by a protocol is defined as the number of bits/s sent on average by each node for route maintenance (building and releasing routes when the sink moves) and for gathering information needed by the sink for deciding whether to move or not and where.
- *Delivery ratio*. The percentage of packets generated at the sensor nodes that are successfully delivered to the sink. In the low traffic scenarios considered here we observe that we are always able to successfully transmit all the packets.
- *Sink sojourn times at the different sites*. This metric captures the time spent by the sink on average at the different sites and is a key figure to correlate sink behavior with network performance (overhead, latency, and energy consumption).

We also examined the mobility pattern followed by the sink for a given sink mobility scheme over different runs to identify common patterns and obtain an in-depth understanding of the rationale behind the sink movements.

6.1 Simulation scenarios and parameters

Our evaluation has been performed according to three major sets of experiments. We have initially focused on a simple scenario in which $n = 400$ sensor nodes generating data at the constant rate of 0.5 bit/s (i.e., a packet is injected into the network around every 13 min) are deployed in a grid-like topology over a square area of side $L = 400$ m. The sensor nodes transmission range R is fixed and equal to 25 m. This means that all nodes which are not on the perimeter of the area have a "cross-like visibility" of their neighbors (i.e., they have four neighbors). A single sink moves through $|S|$ possible sites. Sink sites are arranged into a 4×4, 6×6, 8×8 grid. Data are delivered to the sink according to the routing protocol presented in [7]. The route construction process is sink initiated. The sink floods a packet via which the nodes calculate their hop distance from the sink. Forwarding happens based on this simple (and unchanging) information: A node that is i hops away from the sink will send data packets to one of its neighbors whose distance is $i - 1$. The specific neighbors can be different each time, and one is chosen among the neighbors randomly and uniformly.

Channel capacity and MAC settings are typical of sensor networking (250 Kbps and CSMA/CA, respectively) and we consider the sensor nodes to be all alike. Sensor nodes initial energy is set to 50 Joules. The energy model used in our experiments follows the specifications of the TR 1000 radio transceiver from RF Monolithics [1], i.e., the energy consumption corresponding to transmission, reception, or sleep mode is 14.8 mW, 12.5 mW and 0.016 mW, respectively.[1]

[1] In order to assess the advantages of sink mobility independently of the particular awake/asleep schedule used for energy conservation we did not consider the energy consumed by a node while being idle. This choice corresponds to using an ideal awake/asleep schedule where nodes are awake only when they are involved in message transmission and reception.

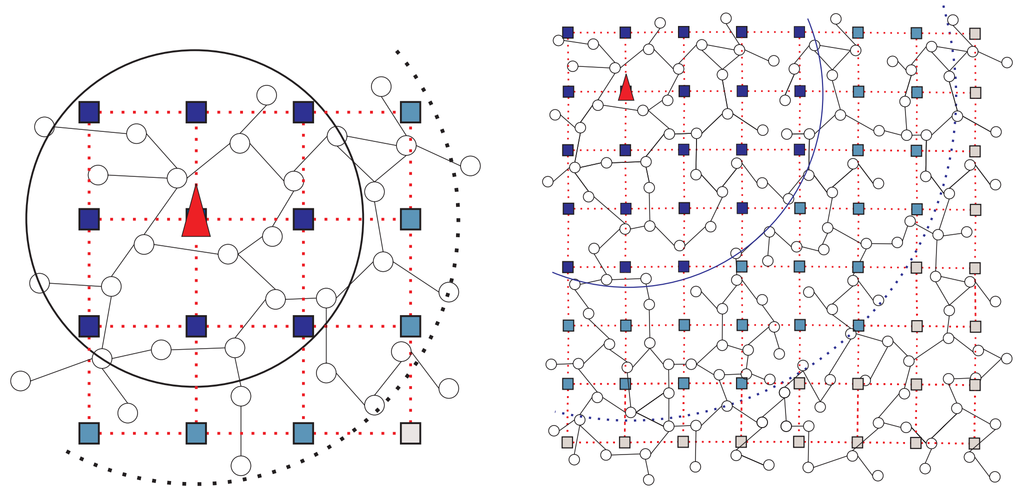

(a) 4 × 4 case (b) 8 × 8 case

Fig. 4 Neighboring sink sites

In this basic scenario the model parameter d_{MAX} has been set to 190 m. The set of sites the sink can move to from its current position depends on d_{MAX} and on the cardinality of S. Figure 4 shows the set of adjacent sink sites when the sink can select where to sojourn among 4×4 and 8×8 sites. The sink is indicated as a triangle. The first circle encloses adjacent sink sites when $d_{\mathrm{MAX}} = 190$ m. The parameter $t_{\min}$ has been varied between 50,000 s and 1,000,000 s.

This first set of experiments aims at demonstrating the bare effectiveness of the proposed MILP and heuristic solutions for extending network lifetime with respect to the static case.

Our second batch of simulations aims at quantifying the impact of key protocol parameters on OPT, GMRE and RM performance. First of all, we have considered a wider nodal transmission range, 30 m. In this case nodes have an "asterisk"-like connectivity (i.e., each non-border node has eight neighbors), which leads to different routes with respect to when the transmission range is 25 m and hence to different values of both c_{ik} and f_{ik}. This allows us to investigate the impact of different forwarding schemes on sink routes and network lifetime.

We have then evaluated the effect of varying d_{MAX}. This changes the set of adjacent sites the sink can move to from its current one. (See Fig. 4, where the larger dotted circle refers to $d_{\mathrm{MAX}} = 325$ m. Values of d_{MAX} greater than 570 m corresponds to unlimited sink mobility.) When $R = 25$ m we have compared the performance of the different mobility schemes for $d_{\mathrm{MAX}} = 190$ m, 325 m, and 1,000 m.

In our third and final set of experiments we assess whether relaxing the assumptions made for node deployment, routing and sink site locations changes the insights gained on the effectiveness of the different sink mobility schemes. In this batch of experiments we have first investigated the case when the sink cannot freely roam throughout the network, but it is constrained to travel only within certain areas. The impact of different data routing has been evaluated by considering a protocol based on geographic forwarding [58, 59] instead of the shortest-paths based routing used before. Finally, we have changed the way sensor nodes are positioned in the deployment area. We consider scattering the nodes randomly and uniformly rather then placing them more regularly on a grid.

Table 1 sums up the wide variety of parameters used in our investigation.

All the presented results achieve a 95% confidence level within a 5% precision.

6.2 First set of experiments: A basic testing scenario

In this section we summarize the results of our first group of experiments, aiming at assessing the performance of the different sink mobility schemes for the basic scenario.

Table 1 Simulation parameters

Parameter	Value
Area side L	300 m, 400 m
Number of sensor nodes	400, 600
Node deployment	grid, random and uniform
Nodal transmission radius	25 m, 30 m
Nodal data rate	0.5 bit/s
Channel data rate	250 Kbs
Data routing	shortest-path like [7], GeRaF [59]
Nodes initial energy	50 J
Energy model	RF Monolithics nodes [1]
Sink sites	16, 36, 64
d_{MAX}	142 m, 190 m, 325 m, 1000 m
$t_{\min}$	50,000 s, . . . , 1,000,000 s

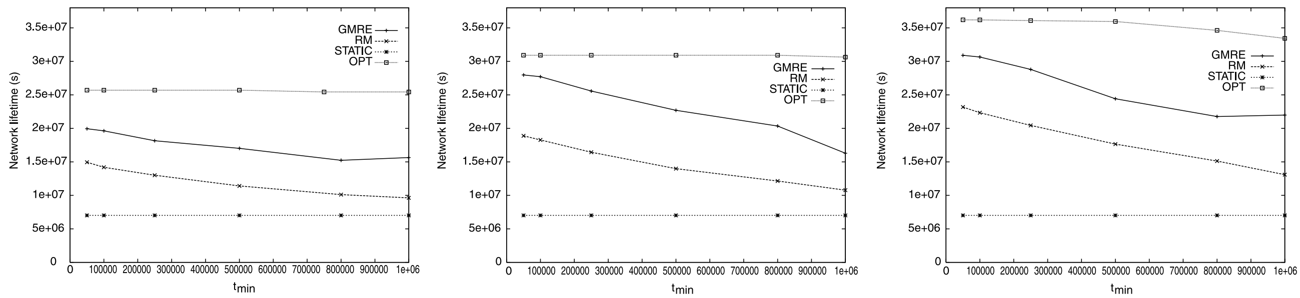

(a) 4 × 4 sink sites (b) 6 × 6 sink sites (c) 8 × 8 sink sites

Fig. 5 Average network lifetimes for increasing t_{min}s

Figure 5 shows the sensor network lifetime in networks with different numbers of sink sites and for values of t_{min} that vary between 50,000 s and 1,000,000 s.

Each figure compares the lifetime obtained by OPT, GMRE, RM and STATIC. The lifetime in the static case is equal to 7,013,801 s (of course independently of t_{min}). This is the time when one of the four sink's neighbors (i.e., the nodes that relay all the network data to the sink) dies because of energy depletion. The other three nodes remain with negligible amount of energy, dying immediately after the first. This leaves the sink isolated from the rest of the network.

The uneven energy consumption of STATIC is shown in Fig. 6, where we depict the distribution of the sensor nodes average residual energy at network lifetime (expressed as a percentage of the initial energy.) The remarkably high variance among the residual energies is due to the different distance of each node from the sink and, in general, to the different number of sensor-to-sink paths to which a node belongs, which implies different number of packets to relay. Nodes along the "cross" centered at the sink tend to be the preferred data relays. The closer these nodes are to the sink, the higher the number of packets they receive and transmit, and consequently the higher their energy consumption. This implies that these are the nodes with the lower residual energy at network lifetime.

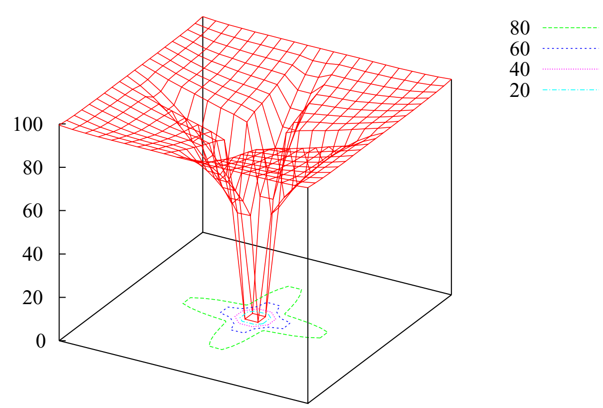

Fig. 6 Average nodal residual energy at lifetime (STATIC, 16 sink sites)

In particular, when the first node dies the other sink's neighbors have a very low residual energy ($\leq$0.03%), the energy at the nodes along the cross arms averages at 71.07%, while 42.75% of the network nodes have more than 95% of their initial energy available! This incapacity of balancing node energy consumption results in short network lifetime and inefficient use of available resources. The sink is soon disconnected from the network, while a large number of the deployed nodes are still fully operational.

The idea of moving the sink stems from the attempt of obviating the sink's neighborhood problem demonstrated by the experiments above. If the sink can move, then the nodes which consume the largest amount of energy for data relaying vary over time. Therefore, energy depletion is more balanced among the nodes, which in turn results in increased network lifetime. Improvements with respect to the static case can be as high as 200% (350%) when the sink moves according to GMRE in scenarios with 16 (64) sink sites. In this case the GMRE lifetime is only 16% (28%) shorter than the OPT lifetime when t_{min} is kept below $\leq 250K$ s.

Improvements on network lifetime are obtained even when the sink moves randomly (RM heuristic). We have observed improvements short of 100% in case the sink can travel to 16 different sites, while longer lifetimes (up to 220% of the STATIC lifetime) are obtained in scenarios with 64 sink sites.

For all the different mobility schemes a higher number of sink sites results in higher lifetimes. The greater number of sites allows the sink to drain the energy of nodes in areas which otherwise it could not visit.

In general, for both GMRE and RM, the lower the t_{min}, the higher the network lifetime. This is due to the fact that at higher t_{min}s the sink cannot finely decide the time to spend at each site, which implies a less uniform energy consumption at the nodes. For very high t_{min}s it is not even possible for the sink to sojourn at all network sites: Lifetime is reached before the sink can visit them all.

Even in the case of OPT mobility, lower t_{min}s correspond to longer lifetimes. In this case, however, the decrease in the network lifetime when t_{min} grows is not as evident as for GMRE and RM. It is interesting to point out the reasons for which this is the case. First of all, in the OPT case $t_{\min}$ is simply a lower bound on the sojourn time: The sink has to stay there for that time, but does not have to move after it, and can stay an (optimum) amount of time after $t_{\min}$ and then move. In case of GMRE and RM dictated sink mobility, the decision about whether to move or not is due every $t_{\min}$ (after which the sink often moves). This has a twofold consequence. On one side, increasing $t_{\min}$ is more imposing for GMRE and RM than for OPT in terms of fine tuning the sojourn time at a site. Moreover, being able of deciding optimum sojourn times implies much lower sink mobility in the OPT case, which corresponds to lower overhead for route management and hence to lower energy consumption with respect to that incurred by GMRE and RM. Secondly, GMRE and RM do not have a global view of the network topology and do not know the network traffic, i.e., how the nodes energy consumption evolves over time, resulting in decreased performance with respect to OPT. The RM heuristic does not enforce any energy-based criterion for sink movement, resulting in the worst performance among all the mobility schemes. Even if GMRE takes into account the nodes residual energy, the decision about whether to move or not, and where, is based on the current status of the network and on a local view of the residual energy. According to the best "greedy" tradition, this could lead to a bad move with respect to global network lifetime maximization. The impact of such bad move is clearly higher for high $t_{\min}$s: The wrong toll is paid for a longer time.

The OPT mobility performance also degrades for higher values of $t_{\min}$. In this case it converges to values that are typical of when the sink is kept static. For instance, for values of $t_{\min}$ approaching 7,013,801 s the MILP model positions the sink at the center of the deployment area and leaves it there (static). However, as explained above, increasing values of $t_{\min}$ are less critical in the case of OPT mobility than in the case of GMRE and RM, and OPT network lifetime values start to decrease steeply at very high $t_{\min}$s (not shown in the figures). Considering that OPT needs global information for deriving optimum sink mobility and sojourn times, and considering also the more "philosophical," algorithmic differences between OPT and GMRE mobility, the fact that our heuristic never obtains network lifetimes more than 28% lower than OPT's, for relatively low sink mobility rate (small $t_{\min}$s), can really be considered an excellent result. It is also worth noticing that GMRE leads to considerably better performance over RM not only in terms of average network lifetime but also in terms of the network lifetime variance. RM performance greatly varies depending on specific sink random movements, resulting in a non-negligible probability of very poor performance. In case the sink can choose among 16 sites and $t_{\min} = 50{,}000$ s ($t_{\min} = 500{,}000$ s), the RM network lifetimes over multiple runs are (almost) uniformly distributed in the range [12M, 17M]s ([6M, 16M]s). In the same scenario GMRE lifetimes vary in the smaller range [19.5M, 20.5M]s ([16.5M, 17.5M]s). This makes GMRE performance much more predictable and much of a better choice over uncontrollable mobility.

Network lifetime betterment due to sink mobility is paid by increased data latency. The reasons are quite clear. First of all, packets that are newly generated while the sink is moving and those in transit toward the sink have to wait until routes to the new position of the sink are established. Furthermore, in order to balance energy depletion, the sink will spend time not only at the center of the deployment area but also along borders. This imposes longer average routes and hence a higher packet latency than the latency experienced when the sink is statically placed at the center. The latter is actually the dominant reason for increased latency in low sink mobility scenarios.

To better understand the average end-to-end latency experienced by a data packet, we investigate the sojourn times of the sink in different parts of the deployment area according to the considered mobility schemes. We observe that the two sink mobility schemes that achieve the highest network lifetimes, i.e., OPT and GMRE, tend to make the sink sojourning at sites on the corners and along the perimeter. This depends on the energy consumption at the nodes when the sink stays at different sites. More precisely, when the sink sojourns at a corner (say, in the lower left part of the deployment area) the highest energy consumption happens close to the sink site, for nodes at the lower and left sides of the area. When the sink is on the perimeter (e.g., on the lower side) the highest energy consumption occurs on the perpendicular line intersecting the lower side at the sink location, and less evidently on the lower side itself. When the sink is located close to the center of the deployment area nodes along the cross centered at the sink site are the most stressed in terms of energy consumption (Fig. 7(a)). In the latter case as there are more nodes within transmission range from the sink they better balance the energy consumption for delivering to the sink the packets generated by the other nodes. This translates into a lower energy consumption experienced at the nodes close to the sink. However, nodes in the central areas always consume energy, independently of where the sink is located. The energy consumption for central nodes can be very high not only when the sink is located at the center but also when it is on the perimeter. Locating the sink at the center thus drains energy from critical nodes whose energy will be continuously depleted throughout the network lifetime. Locating the sink at the corners for long times, instead, appears to be a very promising strategy, as it depletes the energy of nodes which experience very little energy consumption when the sink is at any of the other sites. This motivates the behavior of the GMRE and OPT schemes. OPT leads to further improvements over GMRE

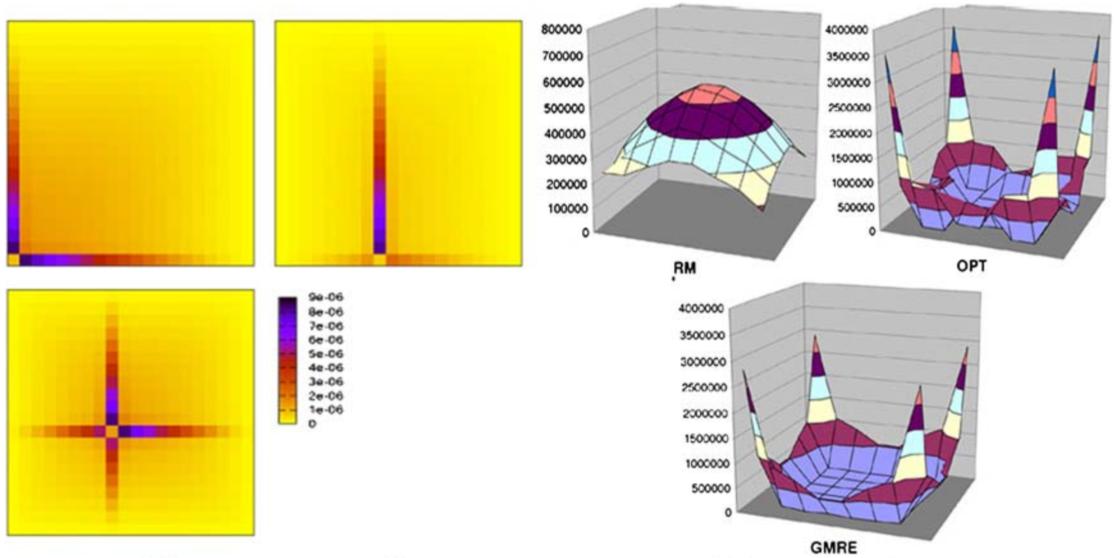

(a) Average energy consumption

(b) Average sojourn times

Fig. 7 Node energy consumption and sink sojourn times

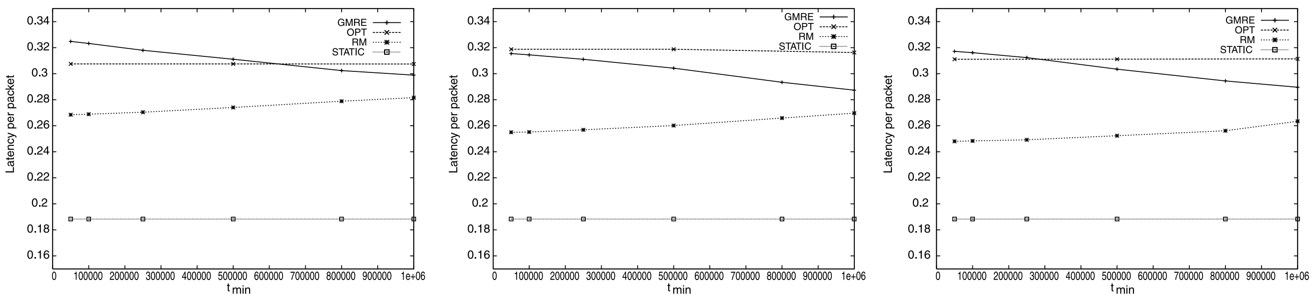

(a) 4×4 sink sites

(b) 6×6 sink sites

(c) 8×8 sink sites

Fig. 8 Average data latency for increasing t_{min}s

as it exploits the available global information on the energy consumption per node when the sink stays at a site and on the traffic to better select the sink sojourn times. For example, if placing the sink at two sites stresses the same nodes, the OPT scheme will tend to avoid spending long sojourn times at both sites. Overall, the finer tuning of the sojourn times leads OPT to better balance the energy consumption among the nodes, and hence to longer lifetimes. At network lifetime around 20% (60%) of the nodes have consumed more than 80% of their initial energy in GMRE (OPT). This was the case for only the 1% of the nodes in STATIC.

Given its stochastic nature, the RM heuristic, positions the sink mostly at the center of the deployment area, resulting in worse load balancing and lower network lifetime. This is clearly shown in Fig. 7(b) which depicts the average sojourn times per site in the case of networks with 64 sites, for OPT, GMRE and RM.

It is now possible to fully understand the latency performance of the different schemes. Figures 8(a) to (c) depict the average latency per packet in OPT, GMRE, RM and STATIC when the number of sink sites varies from 16 to 64. When the sink sojourns at perimeter or at corner sites (which is typical of GMRE and OPT) we know that the lifetime increases. However, these are also the cases when the average length of the routes to the sink increases, which implies, in turn, a higher packet latency. It is thus reasonable to expect that lower latencies are experienced when the sink is statically placed at the center of the sensor deployment area. The RM heuristic, which tends to move the sink to sites located centrally, is the first best mobility scheme in terms of latency. The increase of RM-induced latency with respect to STATIC is expectedly lower when the number of sink sites increases, since in this case central sites are closer to the geographical center of the deployment area. This increase never tops 40%. As t_{min} increases the sink tends to stay less at central sites, leading to higher average latencies experienced by RM packets. The opposite trend is observed for GMRE. For small t_{min}s the sink stays at sites on the corners and on the perimeter, which leads to average latencies up to 30% higher than those experienced in the RM case. When t_{min}

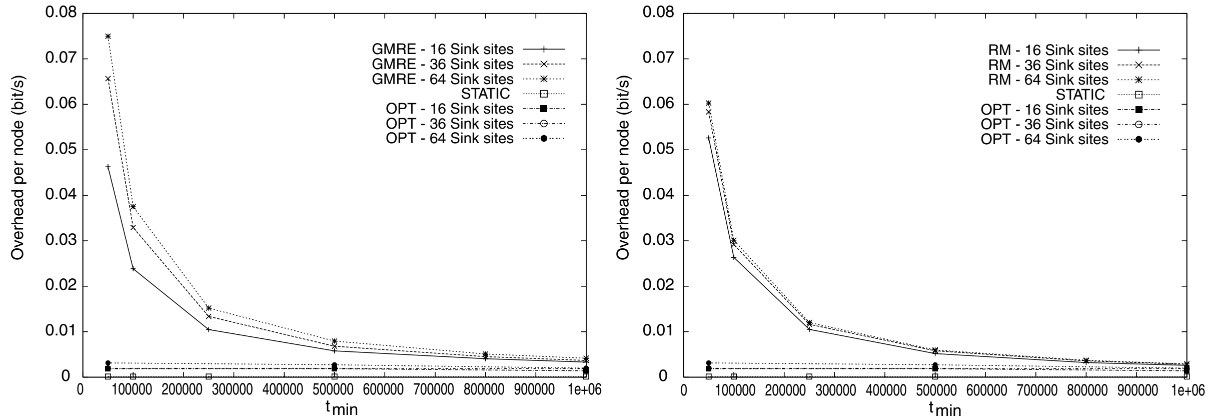

(a) GMRE, OPT and STATIC (b) RM, OPT and STATIC

Fig. 9 Average overhead per node

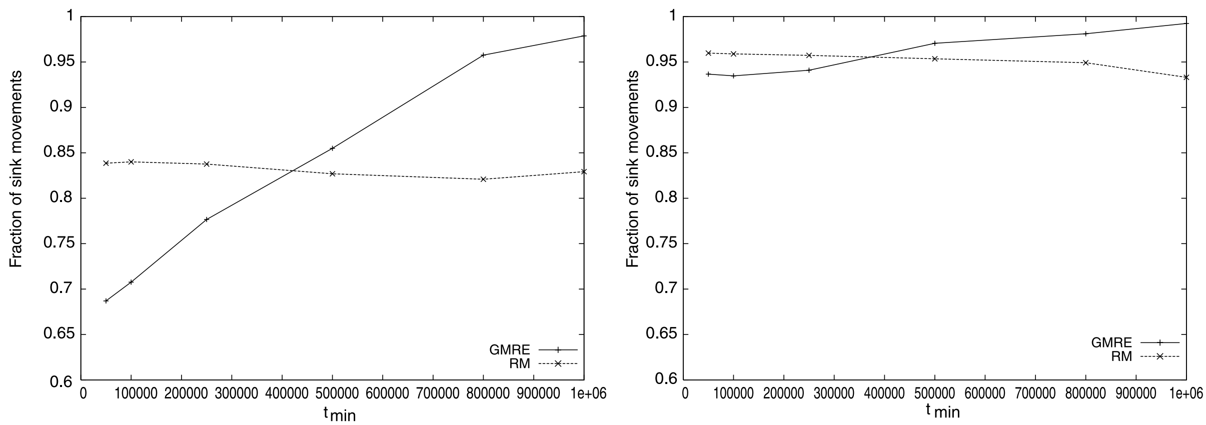

(a) 4×4 sink sites (b) 8×8 sink sites

Fig. 10 Fraction of sink movements

increases the sink sojourns less at corner sites, which implies a decrease in the average latency. Optimum sink mobility is not significantly affected by varying $t_{\min}$ in the selected range, and the latency values are similar to those observed for GMRE.

We have also investigated the overhead of the considered mobility schemes. By overhead we mean the cost incurred by the protocols for route management, i.e., for route set up and release when the sink changes site as well as the cost required for gathering information about residual energy at adjacent sink sites. For OPT we make the ideal assumption that the needed input (c_{ik} and f_{ik}) to the MILP formulation is known. Therefore, the OPT overhead is associated to route maintenance when the sink follows the mobility pattern output by the model. Figure 9 shows the average number of control bits that each node transmits per second (overhead/s) according to OPT, GMRE, RM and STATIC. In particular, Fig. 9(a) shows the overhead/s incurred by GMRE for varying number of sink sites and $t_{\min}$s, and the overhead imposed by OPT and STATIC.

Figure 9(b) depicts the same metric for the RM heuristic and compares it with OPT and STATIC. Both OPT and STATIC impose negligible overhead (they are barely visible in the figure). When the sink is kept static routes need to be computed only once. In the case of OPT mobility, the sink moves from one site i to the next one j only when the whole sojourn time t_i at site i has passed, resulting in a reduced number of movements. According to the GMRE and RM heuristics, instead, the sink makes a decision on whether to move or not every $t_{\min}$. As expected, the higher the $t_{\min}$, the fewer the movements, the lower the overhead per second.

Figures 10(a) and (b) show the fraction of times the sink decides to move according to the two heuristics. (This number is here expressed as the fraction of times the sink actually moves over the times it considers whether to move or not.)

The pictures concern scenarios with 16 and 64 sink sites. We observe that both GMRE and RM move the sink almost always. This justifies the fact that the overhead/s is correlated to the value of $t_{\min}$. In RM the probability of staying at a site

Table 2 Distance (per second) traveled by the sink

	RM			GMRE			OPT	
$t_{\min}$ (sec)	50 K	250 K	500 K	50 K	250 K	500 K	50 K	500 K
Distance traveled (16 sink sites)	0.00263	0.00051	0.00025	0.00181	0.00042	0.000237	0.000077	0.000077
Distance traveled (64 sink sites)	0.00243	0.00048	0.00023	0.00208	0.00043	0.00024	0.0000594	0.0000519

only depends on the number of adjacent sites. The higher the number of sites, the denser they are, and the more likely the sink movements are. This motivates the slight increase in the overhead/s for a higher number of sites. It also justifies the slight decrease in the percentage of sink movements for high $t_{\min}$s. For these values the sink tends to sojourn more in the external parts of the deployment area where the number of adjacent sites of a given site is lower. This results into lower sink mobility. In GMRE, sink mobility depends on whether or not, after staying at least $t_{\min}$s at one site, one of the adjacent sites has a higher residual energy. The higher the $t_{\min}$, the more the nodes close to the current sink site deplete their energy during this sojourn, and the most likely one of the adjacent sites has higher residual energy.

The cost associated with each sink movement is higher for GMRE than for RM. Beside the cost incurred for constructing and releasing routes (a network-wide flooding), GMRE has also to transmit control packets for sentinel identification, for the sink to inquire the sentinels about the residual energy around the adjacent sites, and for delivering this information to the sink. The higher the number of sites, the higher the number of adjacent sites, the higher the cost associated for gathering the energy information, the higher the overhead/s. At $t_{\min} = 50{,}000$ s a 60% increase in GMRE overhead is experienced when the number of sink sites is 64 over the case where only 16 sites are available. However, the increase in overhead/s experienced by GMRE over RM is always quite limited, being the extra cost paid by GMRE for collecting the energy information balanced, at small $t_{\min}$, by the smaller percentage of times the sink moves.

We have finally compared the distance traveled by the sink as dictated by RM, GMRE and OPT. Results are shown in Table 2 (the distance has been normalized to the network lifetime). GMRE and RM have similar performances. As observed in Fig. 10, for small and medium values of $t_{\min}$ RM tends to have the sink moving more frequently, thus resulting in longer distances traveled by the sink. In OPT the sink travels much less. In fact, improvement are observed of up to two orders of magnitude with respect to GMRE and RM.

6.3 Second set of experiments: Impact of varying protocol parameters

In this section we investigate whether varying protocol parameters such as the transmission radius R and d_{MAX} has

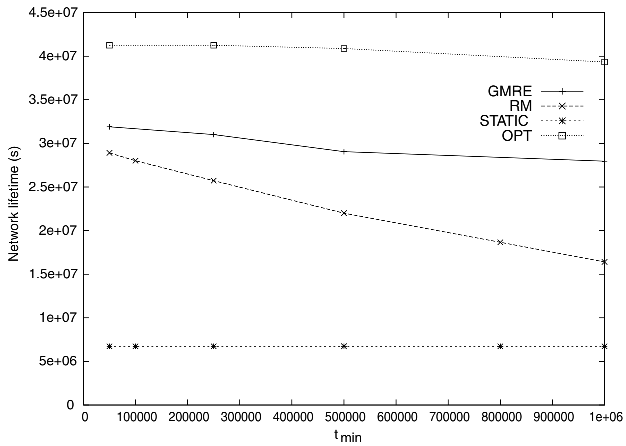

Fig. 11 $R = 30$ m: Average network lifetime

an impact on the relative performance of the sink mobility schemes. By comparing the performance of OPT, GMRE, RM and STATIC when $R = 25$ m (cross-like node visibility) and 30 m (asterisk-like visibility) we are able to understand the impact of different forwarding strategies on optimum sink mobility and on the performance of the two heuristics. By varying d_{MAX} between 190 m and 1,000 m (unbounded sink mobility) we analyze whether looser bounds on the acceptable data latency (i.e., larger d_{MAX} s) result in improved performance or not.

In what follows we show results for networks with 64 sink sites. Experiments with networks with 16 and 36 sites show similar trends.

6.3.1 Varying the transmission range R

Figure 11 depicts the average network lifetime obtained by OPT, GMRE, RM and STATIC when $R = 30$ m. The relative performance of the different schemes shows a trend similar to when $R = 25$ m.

For small $t_{\min}$s OPT achieves an average network lifetime six times longer than STATIC (41,500,000 s). GMRE follows closely, with a lifetime which is only 23% shorter. The RM scheme induces a lifetime which is another 10% shorter. On a wider range of $t_{\min}$s we observe that GMRE lifetime is never lower than 28% of OPT lifetime. RM, instead, obtains a lifetime which is up to 40% lower than GMRE's.

While similar in trend and relative performance, network lifetime values are significantly different from those obtained

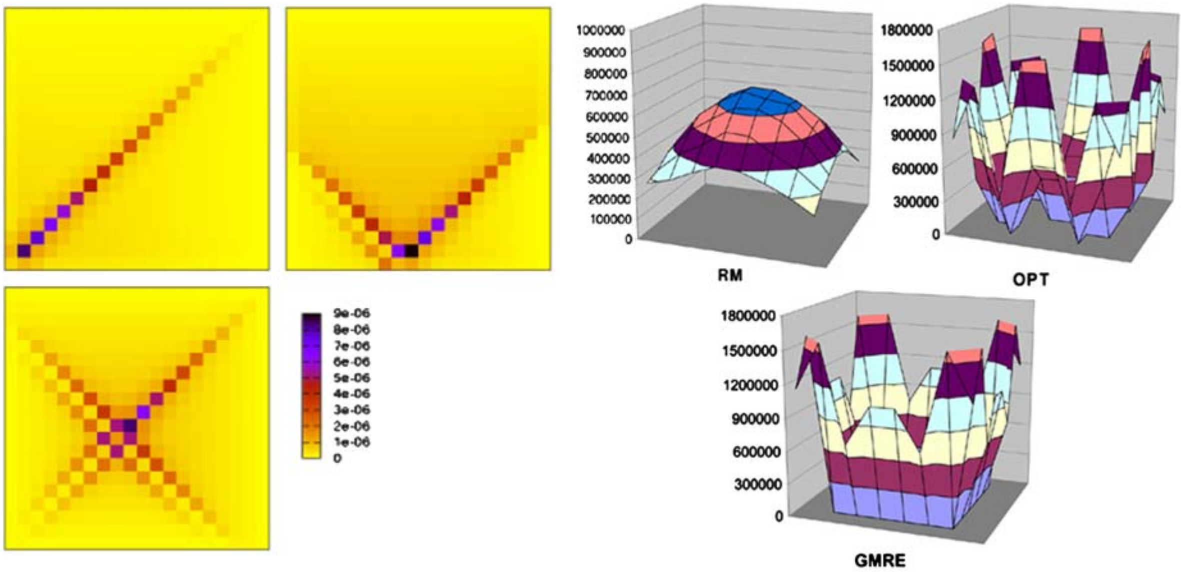

(a) Average energy consumption

(b) Average sojourn times

Fig. 12 $R = 30$ m: Node energy consumption and sink sojourn times

in networks with $R = 25$ m. When $t_{\min} = 50{,}000$ s, for instance, OPT has a lifetime which is 14% higher than when $R = 25$ m. In general, all schemes (but STATIC) experience increased network lifetime when R increases. The improvements are due to the higher nodal density, which induces shorter routes to the sink. This translates into decreased power consumption for data delivery. Shorter routes bring no advantage in the STATIC case, for which we observe no improvement over when $R = 25$ m. The difference here is made by the number of sink's neighbors which are those nodes that relay the packets of every other node, and their own packets. Network lifetime is bound to the death of one of these nodes. Energy consumption at the sink's neighbors depends on the packets they have to relay. The higher the number of sink's neighbors, the lower the number of packets each of them has to forward to the sink. The number of sink's neighbors (four) is the same for the two values of R, resulting in comparable network lifetimes.

For a detailed understanding of the relative performance of the considered sink mobility schemes we have investigated the energy consumption per node when the sink sojourns at different sites k, the sink sojourn times t_k, and the nodes' residual energy at lifetime in the four cases.

The following Fig. 12(a) shows the average energy consumption per second experienced by each node when the sink sojourns at the lower left corner, in the middle of the lower side and at a central site, respectively.

A comparison with Fig. 7(a) immediately shows the different kind of network connectivity, and hence the different data routes obtained when $R = 30$ m. Energy consumption happens now along the diagonal(s) ending at the current sink site. Differently from the case with $R = 25$ m, it appears that corner sites are no longer the most convenient spots, since staying at one of these sites heavily imposes on nodes along the diagonals whose energy is also drained when the sink stays at some of the side and central sites. When $R = 25$ m (Fig. 7(a)), staying at the corners would drain energy from nodes on the sides of the deployment area, which consume negligible energy otherwise.

Figure 12(b) confirms this intuition. According to OPT and GMRE the sink stays mostly on the sides rather than at corner sites. In both cases the sink favors some of the sites on the sides (namely, those close to the corners and those in the middle of the side). This depends on the fact that staying at these sites does not impose on the same nodes, while the other sites on the perimeter drain energy from nodes whose energy is already significantly depleted. We notice that OPT, having a global view of the network and of the energy costs, has the sink sojourning less at the corners, more in the middle of the sides, and, for short times, even at some of the central sites. This allows OPT to perform fine tuning of the sojourn times and hence to obtain more uniform energy consumption network wide. The percentage of nodes with less than 20% (40%) of the residual energy at network lifetime equals 27% (39%) in the OPT case. This percentage decreases to 6% (28%) when the sink moves according to *GMRE*.

Being energy unaware, RM behaves the same as in the case $R = 25$, i.e., the sink mostly sojourns at the central sites. This results into less balanced energy consumption and worse performance in terms of network lifetime.

Figure 13 shows the nodes' residual energy at network lifetime for OPT, RM and GMRE, respectively.

Table 3 Average data latency for different Rs

t_{min}	STATIC		RM		GMRE		OPT	
	50 K	1 M	50 K	1 M	50 K	1 M	50 K	1 M
$R = 25$ m	.188	.188	.248	.263	.317	.289	.311	.311
$R = 30$ m	.134	.134	.175	.182	.214	.214	.198	.197

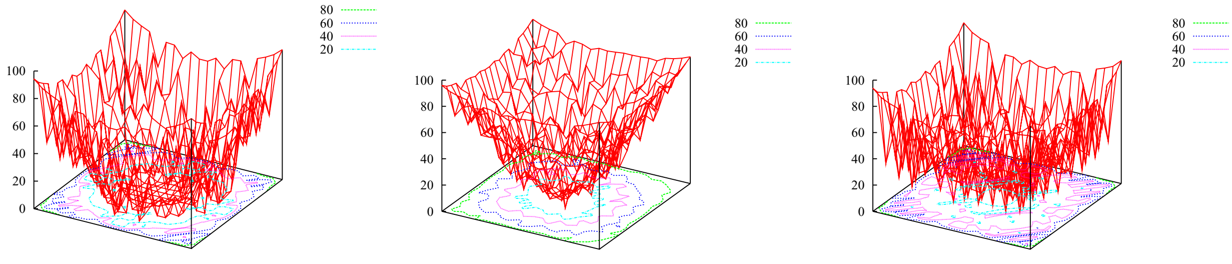

(a) OPT (b) RM (c) GMRE

Fig. 13 $R = 30$ m: Node residual energy at lifetime

The plots refer to best cases, i.e., for each scheme we show results from the one experiment that results in the highest network lifetime. In all the different schemes nodes along the perimeter have high residual energy at network lifetime. This was not the case for $R = 25$ m where the sink sojourns mostly at the corners, imposing on peripheral nodes. When $R = 30$ m nodes along the diagonals to the sink are the ones depleting more energy. Independently of where the sink is, little energy is now drained from perimeter nodes. In case the sink moves according to RM the residual energy is low only at the center of the deployment area (where the sink sojourns most of the time). At the periphery of the deployment area nodes have 60% or more of the initial energy at network lifetime. In GMRE almost all nodes have less than 40% of the initial energy at lifetime. As expected, this improvement over RM is even more evident for OPT. All but the nodes located at the periphery of the deployment area have little energy left at network lifetime.

Now, onto data latency. Table 3 lists the average latency (in seconds) experienced by data packets traveling toward the sink.

By placing and keeping the sink centrally, STATIC always exhibits the best latency performance. RM comes very close for the reason mentioned already (sink stays mostly close to the center). Latency increases when the sink moves according to GMRE and OPT, which have similar performance. As expected, for all schemes the average data packet latency when the transmission range is 30 m decreases with respect to when $R = 25$ m because routes have fewer hops. We observe that RM yields latencies which are up to 21% shorter than GMRE's when $R = 25$ m. When $R = 30$ m RM's average latency is only up to 18% shorter than that imposed by GMRE. In this case GMRE leads the sink more to the sides of the deployment area than to the corners, resulting in shorter routes and improved latency performance. The same is observable for OPT.

The last metric we investigate for the case $R = 30$ m concerns the protocol overhead imposed by the different sink mobility schemes. Not surprisingly, STATIC and OPT have negligible overhead, similar to $R = 25$ m. When $R = 30$ m GMRE experiences increased overhead with respect to the case with $R = 25$ m because of the higher number of nodes which now have to be queried by the sentinels about their residual energy. Furthermore, the increased nodal density imposes a higher number of acknowledgments. These are needed for performing reliable route construction and release. This extra overhead is particularly relevant for smaller t_{min} s. For instance, when $t_{min} = 50{,}000$ s, GMRE experiences an average overhead which is twice as much the bit/s needed when $R = 25$ m. The extra overhead suffered by RM is due to increased flooding complexity, as noticed for GMRE. This leads to an overhead that is almost twofold with respect to when $R = 25$ m.

6.3.2 Varying d_{MAX}

We have assigned the parameter d_{MAX} values in the set $\{190, 325, 1000\}$. Increasing values of d_{MAX} result in increasing numbers of adjacent sink sites. Figure 4 shows the sites the sink can reach from its current site when $d_{MAX} = 190$ m (sites within the smaller circle) and $d_{MAX} = 325$ m (larger dotted circle). When $d_{MAX} = 1{,}000$ m every site is an adjacent site of any other one (i.e., (S, d_{MAX}) defines a complete graph).

Results for the average network lifetime for the three values of d_{MAX} are depicted in Fig. 14(a) ($t_{min} = 50{,}000$ s).

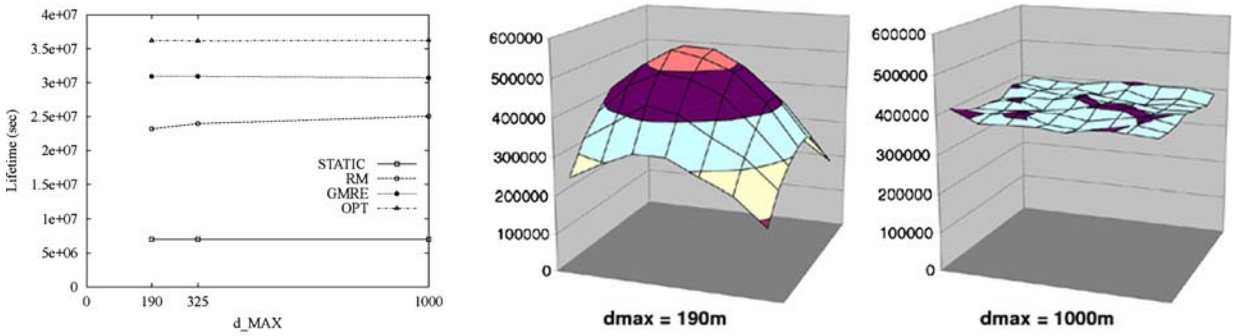

(a) Network lifetime (b) RM sojourn times

Fig. 14 Average network lifetime and RM sojourn time for varying d_{MAX}s

We have considered scenarios with 64 sink sites and $R = 25$ m. (Results for $R = 30$ m show very similar trends.)

Increasing values of d_{MAX} yield slight increases for RM lifetime. This is because the sink is less constrained at staying in the central part of the deployment area (see Fig. 14(b)). As a result, the RM energy consumption among the nodes is more balanced and the lifetime improves. The improvement is however quite limited, never exceeding 8%. For higher t_{min}s (not shown in Fig. 14(a)) network lifetime improves more significantly with increasing d_{MAX}. For high t_{min}s it is critical not to select repeatedly the same sink site or nearby sites that impose on the same nodes. When d_{MAX} increases, the number of adjacent sites increases. It is therefore more unlikely that the sink sojourns for long times in the same area of the network. For the larger values of t_{min} and d_{MAX}, RM achieves a network lifetime which is 19% longer than when $d_{MAX} = 190$ m.

Varying d_{MAX} has no observable impact on the performance of GMRE and OPT. Network lifetime as well as sojourn times at different sites are the same independently of d_{MAX} when $t_{min} \leq 250{,}000$ s. According to both OPT and GMRE when $d_{MAX} \geq 190$ m the sink sojourns only at sites where it is convenient for it to stay. Even if one would expect that increasing the d_{MAX} could be useful in different situations (e.g., when a small d_{MAX} would force the sink to go through sites it would otherwise skip) this is not the case in the considered scenarios. Different d_{MAX} values lead to completely different sink mobility patterns. However, the set of visited sites is the same. More than this: When t_{min} is small, and hence a finer tuning of the sojourn times is possible, OPT and GMRE converge toward the same amount of time spent at the desirable sites independently of d_{MAX}. Therefore, increasing d_{MAX} has only the detrimental effect of increasing data latency, given that the nodes buffer data packets while the sink travels to new sites, now potentially for a longer time. A higher d_{MAX} means the sink can select among a larger set of sites every time it moves. This in turn leads to a higher percentage of sink movements and hence to higher overhead. The larger number of adjacent sites and the higher cost for inquiring about their residual energy makes the increase in overhead particularly evident for GMRE: For unbounded sink movements ($d_{MAX} = 1{,}000$ m) the overhead is three times as much as when $d_{MAX} = 190$ m. This also explains the slight ($< 2\%$) decrease in GMRE network lifetime when $d_{MAX} = 1{,}000$ m and $t_{min} = 50{,}000$ s. For slower sink movement (higher t_{mins}) GMRE can benefit from larger d_{MAXs}. The extra overhead has to be paid less times and at the same time a more global view of the sites residual energy makes wrong moves more unlikely. This explains the 11% increase in GMRE network lifetime when t_{min} is set to 1,000,000 s and d_{MAX} varies from 190 m to 1,000 m.

6.4 Third set of experiments: Challenging initial assumptions

We conclude our investigation about the advantages of sink mobility by discussing whether relaxing the assumptions made in the former sets of experiments affects the performance of our mobility schemes. In particular, we are interested in challenging the decisions made about nodes deployment, data routing, and the freedom of movement of the sink. Our main conclusion here is that, while the values of the metrics of interest (such as network lifetime, latency, etc.) are affected by changing some of the settings (as expected), the relative performance of the mobility schemes does not change. In our experiments, unless specifically noted, the node transmission range R is intended to be 25 m, the number of sink sites is set to 64 and d_{MAX} to 190 m.

We have performed three sets of experiments. The first set concerns limiting the sink movements only to the central sites, i.e., excluding those on the perimeter (where according to the previous experiments, the sink would mostly sojourn). The second batch of experiments gives back to the sink its freedom of roaming throughout the 64 sink sites. What changes in this case is data routing. Data are now delivered to the current position of the sink according to the GeRaF routing protocol [58, 59]. Finally, in the third set of

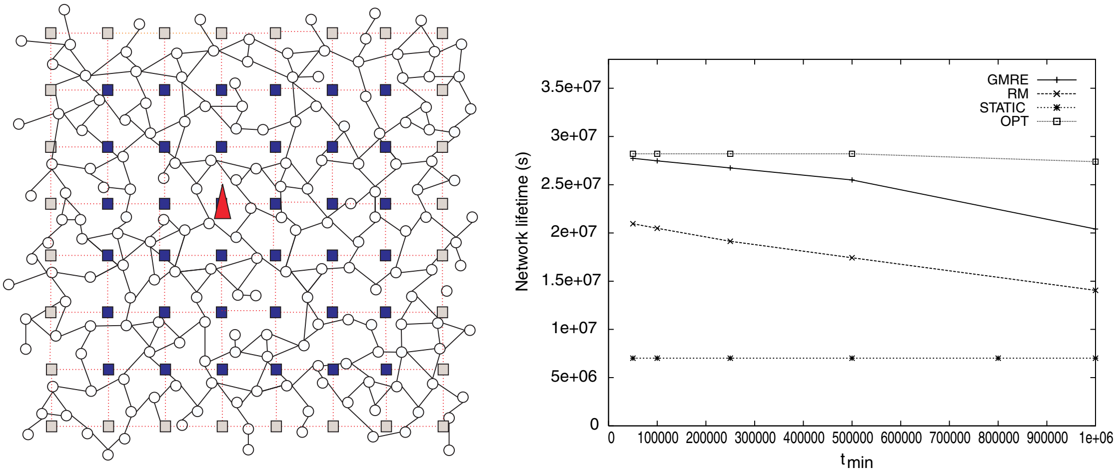

(a) Visitable sink sites (b) Average network lifetime

Fig. 15 Constrained sink mobility: Sink sites and network lifetimes

Fig. 16 Constrained sink mobility: Sink sojourn times ($t_{\min} = 50{,}000$ s)

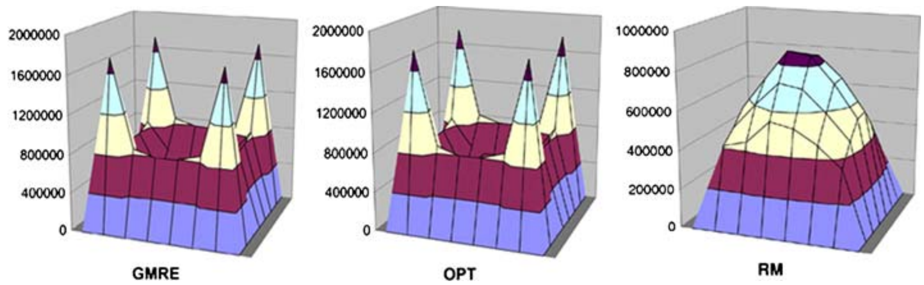

experiments we modify the way sensor nodes are deployed. Nodes are now scattered uniformly and randomly throughout the area rather than placed on a grid. Data are delivered by using the shortest path-like routing as in the basic testing scenario.

6.4.1 Changing sink sites deployment

Figures 15(b) to 18 display results concerning the case of limited mobility of the sink. More precisely, the sink is not allowed to sojourn at any of the 28 sites on the perimeter of the deployment area. The set of sites the sink can sojourn at (called *restricted area*) is displayed in Fig. 15(a) (darker sites).

OPT, GMRE and RM all experience a decrease in network lifetime because energy cannot be drained from nodes at the periphery of the network. However, depending on the specific mobility scheme considered the percentage of lifetime decrease changes quite significantly. It is $\leq 4\%$ for GMRE, $\leq 15\%$ for RM, and $\leq 28\%$ for OPT. A first reason of the greater decrease for OPT is that this scheme is the best in draining energy from *all* the different zones of the deployment area (i.e., it is the more balanced as for energy consumption). Constraining the sink movement impairs OPT's ability of doing it.

A better understanding is achieved by investigating the sink sojourn times and the nodes residual energy over time according to the different schemes. RM sojourn times show that the sink tends to spend the majority of the time at the very center of the deployment area (rightmost plot in Fig. 16). This is more evident in this case rather than in the basic scenarios investigated above. The fact that the sink cannot move to some of the sites corresponds to a lower sink mobility rate, and to the sink sojourning at a restricted group of sink sites. This results in degraded balancing and hence in lower lifetime.

OPT sends the sink preferably to the corners of the restricted area. The sink also spends significant time at the sides of this area, at the four central sites, and at the sites connecting the corners with the central sites. It spends little time in all the other positions. Here is why. When the sink is at the corners, it imposes on nodes which are not too drained when the sink stays at those other sites that are visited for longer times. Specularly, the sites that the sink visits for little time impose high energy consumption on nodes that are already heavily involved in data forwarding while the sink sojourns at those sites it visits longer. OPT exploits its global view of the network for finely selecting the sites to visit and optimally tuning the time spent at each of them. That is why the balance in energy consumption among the nodes achieved

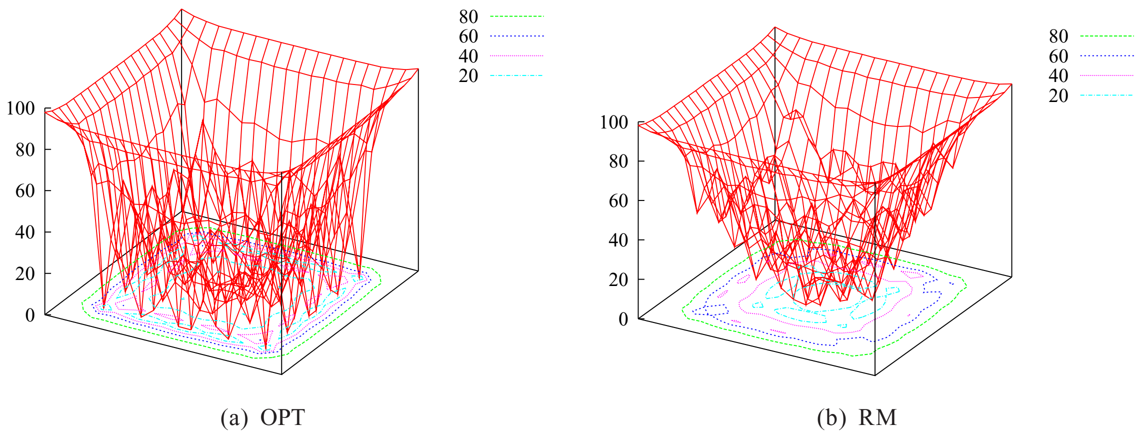

Fig. 17 Constrained sink mobility: Node residual energy at lifetime (OPT and RM)

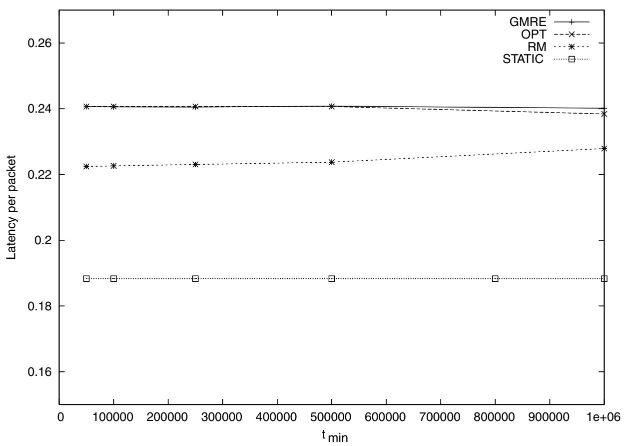

Fig. 18 Constrained sink mobility: Average data latency

by OPT is quite impressive. The node residual energy at the network lifetime is depicted in Fig. 17. Apart from the nodes in the perimeter area that consume little energy having only to transmit their own packets, all other nodes have depleted almost all their initial energy. In particular, of the nodes in the restricted area, 13% (15%, 40%) of the nodes have less than 10% (20%, 40%) of energy left at network lifetime.

GMRE sojourn times for low $t_{\min}$s are depicted in Fig. 16. Quite remarkably, our greedy distributed heuristic is able to perfectly mimic OPT sink mobility, resulting in a network lifetime which is only $\leq 2\%$ lower than OPT's. The local view of the network status is paid with performance degradation at higher $t_{\min}$. At $t_{\min} = 1{,}000{,}000$ s GMRE lifetime is 25% lower than OPT's.

Despite the constraints on sink mobility all three schemes perform well. OPT still obtains fourfold improvement with respect to STATIC (which is of course unaffected by the restrictions). GMRE obtains lifetime improvements that are between threefold and fourfold that of STATIC. RM shows the worst performance. However, it still yields up to 200% improvements over STATIC.

Figure 18 depicts the average data packet latency for OPT, GMRE, RM and STATIC. When the sink is forced to sojourn in the restricted (central) area the average data latency decreases. By comparing these latency values with those of unconstrained sink mobility (Fig. 8(c)) we observe a decrease in latency which is up to 25% for OPT and GMRE and up to 10% for RM. The decrease in all cases is due to the overall shorter routes.

In terms of overhead all schemes but GMRE have basically the same performance as in the unconstrained mobility case. The increase in GMRE overhead is 25%. This is due to the control packet size, which depends on the current position of the sink. If the sink is located in more central sites (as in this case), the number of adjacent sites is larger. The control packet, which contains information about the adjacent sites for sentinel identification purposes, is consequently larger, which explains the higher overhead. Furthermore, a higher number of adjacent sites imposes higher overhead to interrogate the sentinels.

6.4.2 Using geographic instead of shortest paths routing

Table 4 shows the lifetime values for STATIC, RM, GMRE and OPT in case data delivery is performed by GeRaF, the geographically-enabled routing protocol proposed for WSNs in [58, 59]. Nodal transmission radius is set to 25 m, $d_{\mathrm{MAX}} =$ 190 m and the sink can choose where to sojourn among 64 sink sites.

The network lifetime increase achieved by OPT over STATIC is around 300%. GMRE falls short by a mere 20%. RM also improves over STATIC but the improvements can be as low as 10%.

Key to understanding the changes in sink mobility induced by using the GeRaF protocol is the investigation of the energy consumed by the nodes for data forwarding and the sink sojourn times at different sites. Figure 19(a) shows the nodes energy consumption when the sink stays at the lower left

Table 4 GeRaF: Network lifetime

	STATIC		RM		GMRE		OPT	
$t_{\min}$ (sec)	50 K	1 M	50 K	1 M	50 K	1 M	50 K	1 M
Network Lifetime ($\times 10^6$ s)	6.25	6.25	15.6	6.88	20.2	16.5	24.5	21.2

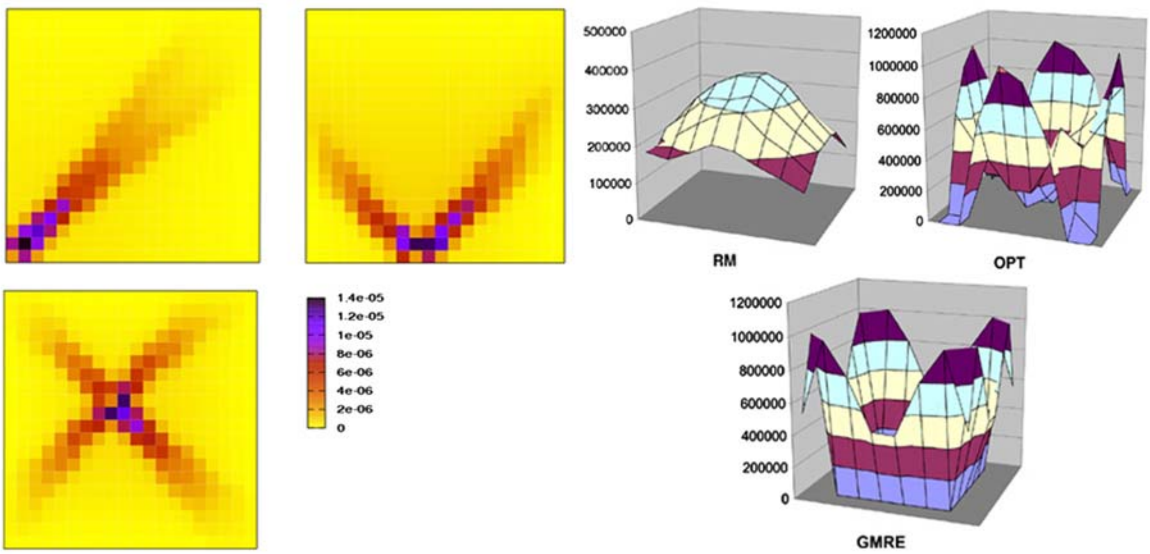

(a) Average energy consumption (b) Average sojourn times

Fig. 19 GeRaF: Node energy consumption and sink sojourn times

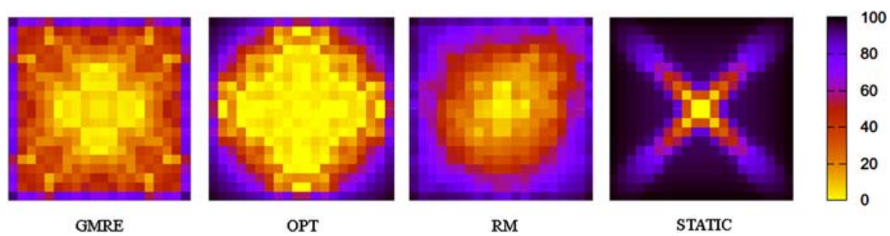

Fig. 20 GeRaF: Node residual energy at lifetime

corner, in the middle of the lower side and at one of the central sites.

The sensor-to-sink routes differ quite significantly from the ones obtained by shortest-path like routing (Fig. 7(a)). This is due to the specific GeRaF forwarding strategy. Whenever a node has a data packet to transmit, it sends it to the one among its neighbors that is closest to the sink itself. This neighbor will relay it further. Therefore, most of the sensor-to-sink routes are made up of nodes that are on "diagonal corridors" terminating at the current sink site. These are the nodes that suffer the highest energy depletion. Similarly to, and more significantly than for the case when $R = 30$ m and routing is shortest path-like, corner sites are no longer the best option, or at least going to the side and central sites is a viable alternative option (since both corner sites and central/side sites impose on overlapping sets of nodes).

The OPT sojourn times confirm this reasoning (Fig. 19(b)). OPT sends the sink for longer times along the perimeter of the deployment area, and, for shorter times, also in the central area. The corner areas are now totally avoided by the sink, since visiting them would stress on the same nodes already drained when the sink stays at the other sites.

As often happens for the omniscient OPT, energy consumption is remarkably balanced: 21.75% (36.25%) of the nodes have less than 5% (20%) of their initial energy at network lifetime. Figure 20, second plot, shows the OPT nodal residual energy at network lifetime (higher residual energy is depicted by a darker color). A large number of nodes (lighter color) have almost no energy left at network lifetime. A better appreciation of how sink mobility is crucial for more balanced energy consumption is obtained by comparing the OPT residual energy with that of STATIC (Fig. 20, rightmost plot). When the sink is statically placed in the center of the deployment area, nodes on the line from each corner to the sink experience high energy consumption. The closer to the sink, the higher the power required to forward data.

At network lifetime, only the sink neighbors have little (or none) residual energy left and only 2% of the nodes have consumed at least 80% of the initial energy. As it is clear from the figure, most of the network nodes have almost all their initial energy.

As in previous experiments, OPT obtains the most uniform distribution of residual energy and corresponding highest lifetime. This is because of its global knowledge of the network topology and data relay costs that enables optimal choice of sink sites and sojourn times.

The GMRE heuristic, being distributed and localized, is able to perform only a coarser "tuning" of the sojourn times. By greedily choosing the best possible next site GMRE can avoid visiting sites whose surrounding nodes have already been drained. Figure 19(b) shows clearly that GMRE recognizes that staying at the corners is the alternative to visiting sites on the sides and at the center. However, lack of global knowledge of key network parameters does not allow GMRE to find the best energy consumption balancing solution. As shown in Table 4 and in Fig. 20 (first plot) GMRE coarser tuning of the sojourn times results into lower network lifetime and more limited parts of the deployment area with little or no energy at network lifetime. When the first node dies 11.5% (32%) of the nodes have less than 5% (20%) of their initial energy.

RM has the sink blindly sojourning at the center of the deployment area, thus stressing the same central nodes. This is why it experiences much worse performance than GMRE and OPT. At network lifetime only 1.5% of the nodes have less than 5% of their initial energy, and only 5% have consumed at least 80% of it (Fig. 20, third plot).

As latency is concerned, RM and OPT have similar performances. In both schemes the sink avoids the most external areas. The end-to-end latency increase with respect to STATIC is never higher than 45%. GMRE pays a further 10 to 18% increase due to the sink spending most of the time at sites in the corners of the deployment area, which impose longer routes. Finally, overhead performances are comparable to those observed in the basic scenario.

6.4.3 Changing node deployment

For the set of experiments concerning varying the deployment of the network nodes we have considered topologies generated by scattering 600 nodes randomly and uniformly in a square area of 300 × 300 square meters. The nodes transmission radius R has been set to 30 m. (Results for the case with $R = 25$ m show a similar trend.) In this scenario, each node has an average of 18 neighboring nodes. Sink sites are arranged according to a 6 × 6 grid. All the other parameter values have been set as in the basic scenario, except for d_{MAX} which is now 142 m. This allows the sink to move to at most 21 possible adjacent sites.

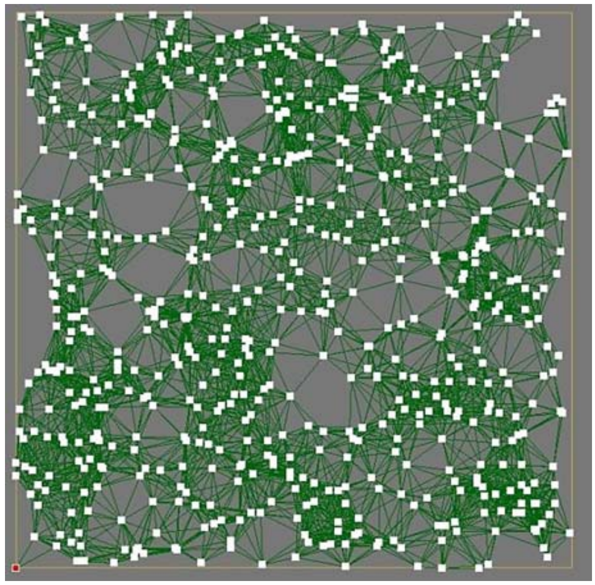

Fig. 21 Random deployment of network nodes

The results we show here refer to the topology displayed in Fig. 21. (Experiments run on 10 other topologies show similar trends.)

In Fig. 22(a) we show the energy consumed by each node, per second, while the sink sojourns at the lower left corner, in the middle of the lower side, and in the center of the deployment area, respectively. (The darker the color, the higher the energy consumption.)

In all cases, the energy consumption pattern induced by shortest paths routing is more "diffuse" as compared to the more regular patterns observed in the grid-based deployment. This makes it more difficult to find a combination of sojourn times which results in balanced energy consumption. We notice that nodes a few hops from the sink site are subject to non negligible energy draining, sometimes comparable to that of the sink's neighbors. Independently of the sink's position, nodes in the central area tend to be among the most stressed. As a consequence, central sink sites, that impose the most on central nodes tend not to be selected by OPT and GMRE. This is clearly shown in Fig. 22(b) which depicts the average sojourn times for the three mobility schemes when $t_{\min} = 50{,}000$ s.

Given that RM is completely energy unaware, it drives the sink as observed in previous experiments. Things change substantially for OPT and GMRE. In the case of optimal sink movements, the choice of the sink sites is carefully performed so to impose energy consumption on almost completely disjoint sets of nodes, a more challenging task in this case! By investigating the nodal energy cost at each node, we observed that OPT leads the sink to sites that with very few exceptions do not impose on the same nodes. In addition, sojourn times are selected so that the most

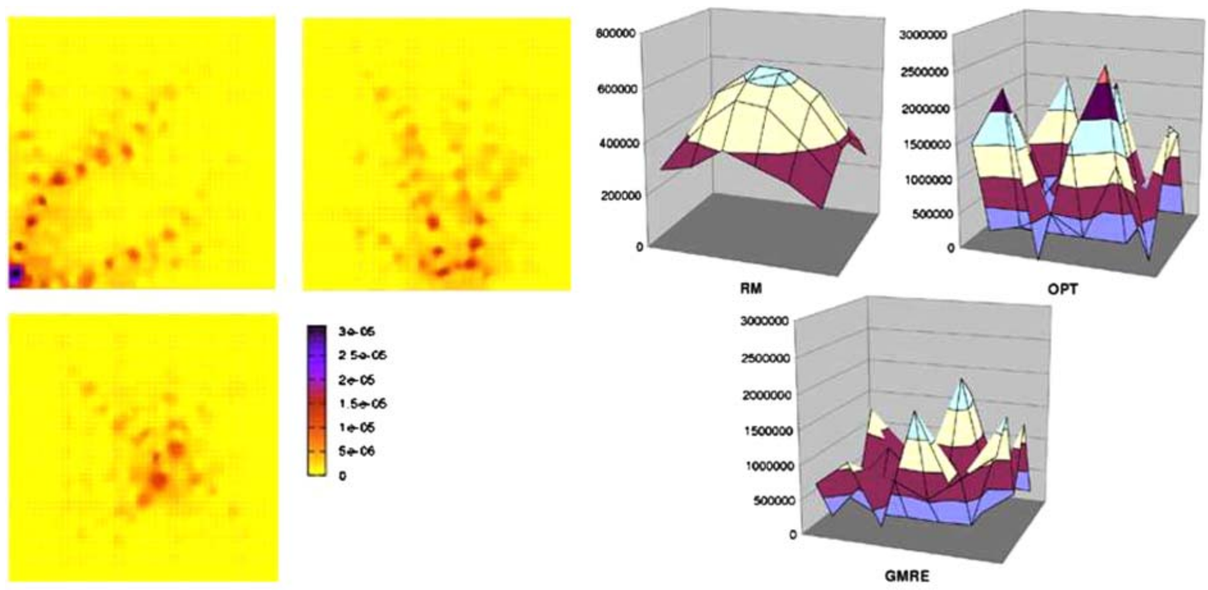

(a) Energy consumption (b) Sojourn times

Fig. 22 Random node deployment: Node energy consumption and sink sojourn times

Fig. 23 Random node deployment: Node residual energy at lifetime

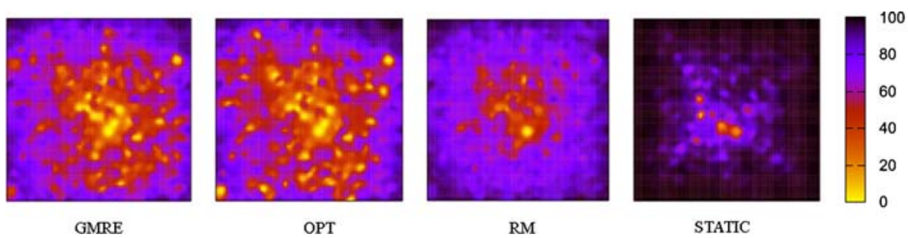

energy-stressed nodes experience comparable drainage, i.e., their residual energy is fairly balanced.

GMRE mimics OPT operations. However, its myopic view of the network does not allow it to optimally select sites and times. This results in the sink staying for long times at sites that sometimes impose on overlapping sets of nodes. We observed that these sets of common nodes, although relatively small, are bigger than those of OPT, which justifies GMRE (limited) performance degradation in terms of network lifetime and of balancing energy consumption.

The results shown refer to an optimized version of GMRE that takes into account the residual energy of all the critical nodes (not only the sink neighbors) for deciding about sink movements. A critical node is a node that experiences a very high energy consumption (comparable to that of a sink neighbor) when the sink stays at a given site. According to this variant, the sentinels learn who are the critical nodes around them by collecting information about the estimated costs at a node when the sink is at their associated site. By achieving a more and more accurate estimate of the set of critical nodes, and by inquiring them on their residual energy and energy consumption, GMRE is able to make an increasingly better guess of which node will have minimum residual energy after the sink has moved to that site for $t_{\min}$. This knowledge is used for allowing the sink to make better decisions about the next adjacent site to move to. Given the incremental nature of critical node discovery (it is a learning process), the process of selecting the next sink site improves in time.

The ability of the four schemes in distributing energy consumption throughout the network is shown in Fig. 23, which displays the nodes residual energy at lifetime for GMRE, OPT, RM and STATIC (from left to right). The plots refer to best cases, i.e., for each scheme we show results from the one experiment that leads to the highest network lifetime. The parameter $t_{\min}$ is set to 50,000 s.

As expected, STATIC shows the worst performance: At network lifetime the overwhelming majority of network nodes (95%) have more than 80% of their initial energy left. Network lifetime is reached because of the complete energy depletion of just a few nodes. Only 0.62% (1.88%) of the nodes remain with less that 20% (60%) of their initial energy. Moving the sink leads to remarkable performance improvement. Even with uncontrolled sink mobility (RM), energy consumption is much more distributed throughout the nodes. Only 45% of the nodes have more than 80% of their initial energy at network lifetime. The 0.62% (15.09%) of network nodes have consumed more than 80% (60%) of

Table 5 Random node deployment: Network lifetime

	STATIC	RM			GMRE			OPT		
$t_{\min}$ (sec)	50 K	50 K	250 K	500 K	50 K	250 K	500 K	50 K	250 K	500 K
Network lifetime ($\times 10^6$ s)	5.2	15	14	13	24	23	20	25.9	25.9	25.8

their initial energy at lifetime. However, RM's marked preference for sending the sink at central sites leads to little energy drainage from the peripheral nodes. Controlling sink mobility (as in GMRE and OPT) is the needed step to leveling energy consumption. In both GMRE and OPT all nodes, but the very external ones (those literally in the deployment area borders), have consumed considerable energy at lifetime. Eighty percent of the nodes have less than 80% of their energy left when the sink moves according to GMRE. In this case, about 39% of the nodes have actually consumed more than 40% of their energy at lifetime. We obtain similar values for OPT mandated mobility. However, in this case, given the better capacity of planning sink movement, there are more nodes with very little energy left. The 7.54% of OPT nodes is left with less than 20% of their initial energy, vs. the 5% of GMRE nodes. This justifies OPT's slightly higher network lifetime.

Table 5 reports the network lifetimes obtained by STATIC, RM, GMRE and OPT when $t_{\min} = 50{,}000$ s, 250,000 s, and 500,000 s. With respect to STATIC, OPT achieves a fivefold improvement. For small $t_{\min}$s GMRE performance is very close to OPT's, with lifetime decreases always below 8%. When $t_{\min}$ increases, the GMRE learning process is not able to converge fast enough to good estimates of the critical nodes, leading to an increasing probability of "bad moves." These moves impose an energy toll which is particularly detrimental to network performance at higher $t_{\min}$s. However, GMRE lifetime decreases over OPT's are never higher than 20%. The energy-unawareness is paid by RM by exhibiting much worse performance. For all $t_{\min}$s RM network lifetime is, on average, from 40 to 50% shorter than GMRE's. In addition, RM lifetime is also highly variable. RM obtains alternatively poor and good performance results depending on the particular run and on the specific random movements followed by the sink. When $t_{\min} = 500{,}000$ s or 1,000,000 s RM network lifetime can be as low as 1,600,000 s, while it is never lower than about 20,000,000 s for GMRE. When $t_{\min} = 250{,}000$ s (50,000 s) RM achieves network lifetimes as low as 7,000,000 s (12,000,000 s), which pales in comparison to GMRE lifetime, which is never lower than 22,600,000 s (23,400,000 s).

Results concerning data latency are consistent with those observed in previous scenarios. STATIC and RM, keeping the sink at central sites, achieve lower latencies because of shorter routes. At $t_{\min} = 50{,}000$ s STATIC imposes an average latency of 0.1 s. Average latency in RM is only 30% higher. OPT's latency is close to GMRE's. Since the sink mostly stays at peripheral sink sites, routes to the sink are on average longer, and hence data delivery time from sensor to sink is higher than for STATIC and RM. However, for small $t_{\min}$s the average latency of the two schemes is only 11% higher than the one imposed by RM.

7 Conclusions

This paper is concerned with prolonging the lifetime of wireless sensor networks. To achieve this goal, we exploit the mobility of the network sink so that, by sojourning in the vicinity of different sensors, energy consumption is more uniformly distributed throughout the nodes. As a consequence, nodes lifetime, and the network lifetime are increased. We have introduced three schemes that represent different solutions for sink mobility. The first scheme, termed OPT, computes optimal sink routes and sojourn times based on a new MILP formulation that considers realistic parameters of wireless sensor networking and sink mobility. This scheme achieves the best possible network lifetime by deciding sink movements based on nodal transmission costs in a centralized way that is typical of LP models. Controlled sink mobility is also generated by our distributed heuristic GMRE. In this case the sink greedily travels toward those areas whose nodes have the highest residual energy. For the sake of benchmarking we have also introduced RM, a simple distributed scheme according to which the sink moves randomly—hence uncontrollably—among the nodes. Via ns2-based simulations on a host of different scenarios, we have demonstrated that whether controlled or not, mobility is key for improving network lifetime. In particular, controlled mobility is effective in prolonging lifetime up to 6 times than when the sink does not move.

Acknowledgments This work was supported in part by the European FP6 027227 IP Project "E-Sense (Capturing Ambient Intelligence for Mobile Communications through Wireless Sensor Networks)."

References

1. ASH transceiver designer's guide (May 19, 2004). www.rfm.com.
2. K. Akkaya and M. Younis, Energy-aware to mobile gateway in wireless sensor networks, in: *Proceedings of the IEEE Globecom 2004 Workshops*, Dallas (TX) (November 29–December 3, 2004), pp. 16–21.

3. I.F. Akyildiz, W. Su, Y. Sanakarasubramaniam and E. Cayirci, Wireless sensor networks: A survey, Computer Networks 38(4) (2002) 393–422.
4. P. Baruah, R. Urgaonkar and B. Krishnamachari, Learning-enforced time domain routing to mobile sinks in wireless sensor fields, in: *Proceeding of the 29th Annual IEEE International Conference on Local Computer Networks, LCN 2004*, Tampa FL (November 16–18, 2004), pp. 525–532.
5. S. Basagni, A. Carosi, E. Melachrinoudis, C. Petrioli and M.Z. Wang, A new MILP formulation and distributed protocols for wireless sensor networks lifetime maximization, in: *Proceedings of the IEEE International Conference on Communications, ICC 2006*, Istanbul Turkey (June 11–15, 2006).
6. S. Basagni, A. Carosi, E. Melachrinoudis, C. Petrioli and M.Z. Wang, Protocols and model for sink mobility in wireless sensor networks, ACM Mobile Computing and Communication Review, MC^2R 10(4) (October 2006) 28–30.
7. S. Basagni, M. Elia and R. Ghosh, ViBES: Virtual backbone for energy saving in wireless sensor networks, in: *Proceedings of the IEEE Military Communication Conference, MILCOM 2004*, Monterey CA (October 31–November 3, 2004).
8. E.H. Callaway, Jr., *Wireless Sensor Networks: Architectures and Protocols*. Auerbach Publications, Boca Raton FL (August 2003).
9. A. Chakrabarti, A. Sabharwal and B. Aazhang, Using predictable observer mobility for power efficient design of sensor networks. In: F. Zhao and L. Guibas (eds.), *Proceedings of the Second International Workshop on Information Processing in Sensor Networks, IPSN 2003*, LNCS 2634, Palo Alto CA (April 22–23, 2003), pp. 129–145.
10. J.-H. Chang and L. Tassiulas, Maximum lifetime routing in wireless sensor networks, IEEE Transactions on Networking 12(4) (2004) 609–619.
11. I. Chatzigiannakis and S. Nikoletseas, An adaptive compulsory protocol for basic communications in highly changing ad-hoc mobile networks, in: *Proceedings of the International Parallel and Distributed Processing Symposium, IPDPS 2002*, Fort Lauderdale FL (April 15–19, 2002), pp. 193–202.
12. I. Chatzigiannakis, S. Nikoletseas and P. Spirakis, An efficient communication strategy for ad-hoc mobile networks, in: *Proceedings of the Twentieth Annual ACM Symposium on Principles of Distributed Computing, PODC 2001*, Newport, RI (August 26–29, 2001), pp. 320–332.
13. E. Ekici, Y. Gu and D. Bozdag, Mobility-based communication in wireless sensor networks, IEEE Communications Magazine 44(7) (2006) 56–62.
14. S.R. Gandham, M. Dawande, R. Prakash and S. Venkatesan, Energy efficient schemes for wireless sensor networks with multiple mobile base stations, in: *Proceedings of IEEE Globecom 2003*, vol. 1, San Francisco, CA (December 1–5, 2003), pp. 377–381.
15. M.R. Garey and D.S. Johnson, *Computers and Intractability. A Guide to the Theory of NP-Completeness*. W.H. Freeman and Co., New York, NY (1979).
16. D.K. Goldenberg, J. Lin, A.S. Morse, B.E. Rosen and Y.R. Yang, Towards mobility as a network control primitive, in: *Proceedings of the 5th ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc 2004*, Roppongi Hills, Tokyo, Japan (May 24–26, 2004), pp. 163–174.
17. M. Grossglauser and D.N.C. Tse, Mobility increases the capacity of ad-hoc wireless networks, IEEE/ACM Transactions on Networking 10(4) (2002) 477–486.
18. A. Howard, M.J. Matarić and G.S. Sukhatme, An incremental deployment algorithm for mobile sensor networks, in: G.S. Sukhatme (ed.), Kluwer Autonomous Robots, Special Issue on Intelligent Embedded Systems 13(2) (2002) 113–126.
19. A. Howard, M.J. Matarić and G.S. Sukhatme, Mobile sensor networks deployment using potential fields: A distributed, scalable solution to the area coverage problem, in: *Proceedings of the 6th International Symposium on Distributed Autonomous Robotic Systems, DARS 2002*, Fukuoka Japan (June 25–27, 2002).
20. X. Hu, Y. Liu, M.J. Lee and T.N. Saadawi, A novel route update design for wireless sensor networks, ACM Mobile Computing and Communications Review 8(1) (2004) 18–26.
21. K. Hwang, J. In, Y. Yun and D. Eom, Dynamic sink oriented tree algorithm for efficient target traking of multiple mobile sink users in wide sensor fields, in: *Proceedings of the 60th IEEE Vehicular Technology Conference, VTC 2004-Fall*, vol. 7, Los Angeles, CA (September 26–29, 2004) pp. 4607–4610.
22. C. Intanagonwiwat, R. Govindan, D. Estrin, J. Heidemann and F. Silva, Directed diffusion for wireless sensor networking, IEEE/ACM Transactions on Networking 11(1) (2003) 2–16.
23. S. Jain, R. C. Shah, W. Brunette, G. Borriello and S. Roy, Exploiting mobility for energy-efficient data collection in sensor networks, in: *Proceedings of the IEEE Workshop on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks, WiOpt 2004*, Cambridge UK (March 24–26, 2004).
24. D. Jea, A.A. Somasundara and M.B. Srivastava, Multiple controlled mobile elements (data mules) for data collection in sensor networks. In: V. K. Prasanna, S. Iyengar, P. G. Spirakis and M. Welsh (eds.), *Proceedings of the 1st IEEE International Conference on Distributed Computing in Sensor Systems, DCOSS 2005*, vol. 3560 of *LNCS*, Marina del Rey, CA (June 30–July 1, 2005), pp. 244–257.
25. J.G. Jetcheva and D.B. Johnson, Adaptive demand-driven multicast routing in multi-hop wireless ad hoc networks, in: *Proceedings of the 2nd ACM International Symposium on Mobile Ad Hoc Networking & Computing, MobiHoc 2001*, Long Beach, CA (October 4–5, 2001), pp. 33–44.
26. P. Juang, H. Oki, Y. Wang, M. Martonosi, L.-S. Peh and D. Rubenstein, Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with zebranet, in: *Proceedings of the 10th International Conference on Architecural Support for Programming Languages and Operating Systems, ASPLOS-X*, San Jose, CA (October 5–9, 2002), pp. 96–107.
27. H. Jun, W. Zhao, M.H. Ammar, E.W. Zegura and C. Lee, Trading latency for energy in wireless ad hoc networks using message ferrying, in: *Proceedings of the 3rd IEEE International Conference on Pervasive Computing and Communications, PerComm 2005*, Kawai, HA (March 8–12, 2005), pp. 220–225.
28. A. Kansal, A.A. Somasundara, D.D. Jea, M.B. Srivastava and D. Estrin, Intelligent fluid infrastructure for embedded networks, in: *Proceedings of the 2nd ACM/SIGMOBILE International Conference on Mobile Systems, Applications, and Services, MobySys 2004*, Boston, MA (June 6–9, 2004), pp. 111–124.
29. H.S. Kim, T.F. Abdelzaher and W.H. Kwon, Minimum energy asynchronous dissemination to mobile sinks in wireless sensor networks, in: *Proceedings of the First International Conference on Embedded Networked Sensor Systems, SenSys 2003*, Los Angeles, CA (November 5–7, 2003), pp. 193–204.
30. Q. Li and D. Rus, Sending messages to mobile users in disconnected ad-hoc wireless networks, in: *Proceedings of 6th ACM Annual International Conference on Mobile Computing and Networking, MobiCom 2000*, Boston, MA (August 6–11, 2000), pp. 44–55.
31. J. Luo and J.-P. Hubaux, Joint mobility and routing for lifetime elongation in wireless sensor networks, in: *Proceedings of IEEE Infocom 2005*, vol. 3, Miami, FL (March 13–17, 2005), pp. 1735–1746.
32. V.P. Mhatre, C. Rosenberg, D. Kofman, R. Mazumdar and N. Shroff, A minimum cost heterogeneous sensor network with a lifetime contraint, IEEE Transactions on Mobile Computing 4(1) (2005) 4–15.

33. C.E. Miller, A.W. Tucker and R.A. Zemlin, Integer programming formulation of traveling salesman problems, Journal of the ACM 7(4) (1960) 311–325.
34. M. Mukarram Bin Tariq, M.H. Ammar and E.W. Zegura, Message ferry route design for sparse ad hoc networks with mobile nodes, in: *Proceedings of the 7th ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc 2006*, Firenze, Italy (May 22–25, 2006), pp. 37–48.
35. I. Papadimitriou and L. Georgiadis, Maximum lifetime routing to mobile sink in wireless sensor networks, in: *Proceedings of the 2005 International Conference on Software, Telecommunications and Computer Networks, SoftCOM 2005*, Split, Croatia (September 15–17, 2005).
36. R. Rao and G. Kesidis, Purposeful mobility for relaying and surveillance in mobile ad hoc sensor networks, IEEE Transactions on Mobile Computing 3(3) (2004) 225–232.
37. A. Savvides and M.B. Srivastava, Location discovery. In: S. Basagni, M. Conti, S. Giordano and I. Stojmenovic (eds.), *Mobile Ad Hoc Networking*, chap. 8, pp. 231–254. IEEE Press and John Wiley and Sons, Inc., Piscataway, NJ and New York, NY, (2004).
38. R.C. Shah, S. Roy, S. Jain and W. Brunette, Data MULEs: Modeling a three-tier architecture for sparse sensor networks, in: *Proceedings of the First IEEE International Workshop on Sensor Network Protocols and Applications, SNPA 2003*, Anchorage, AK (May 11, 2003), pp. 30–41.
39. T. Small and Z.J. Haas, Resource and performance tradeoffs in delay-tolerant wireless networks, in: *Proceedings of the ACM SIGCOMM 2005 Workshop on Delay-Tolerant Networking and Related Networks, WDTN 2005*, Philadelphia, PA (August 22–26, 2005).
40. A.A. Somasundara, A. Ramamoorthy and M.B. Srivastava, Mobile element scheduling for efficient data collection in wireless sensor networks with dynamic deadlines, in: *Proceedings of the 25th IEEE International Real-Time Systems Symposium, RTSS 2004*, Lisbon, Portugal (December 5–8, 2004), pp. 296–305,.
41. L. Song and D. Hatzinakos, Dense wireless sensor networks with mobile sinks, in: *Proceedings of the IEEE International Conference on Acoustic Speech, and Signal Processing, ICASSP 2005*, vol. 3, Philadelphia, PA (March 18–23, 2005), pp. 677–680.
42. The VINT Project. *The ns Manual* (2006). http://www.isi.edu/nsnam/ns/.
43. Y. Tirta, B. Lau, N. Malhotra, S. Bagchi, Z. Li and Y.-H. Lu, IEEE Monograph on Sensor Network Operations, in: S. Phoha and T. F. La Porta (eds.), *chapter Controlled Mobility for Efficient Data Gathering in Sensor Networks with Passively Mobile Nodes*. IEEE Press and John Wiley and Sons, Inc., Piscataway, NJ (December 2, 2005).
44. Y. Tirta, Z. Li, Y.-H. Lu and S. Bagchi, Efficient collection of sensor data in remote fields using mobile collectors, in: *Proceedings of the 13th International Conference on Computer Communications and Networks, ICCCN 2004*, Chicago, IL (October 11–13, 2004), pp. 515–519.
45. L. Tong, Q. Zhao and S. Adireddy, Sensor networks with mobile agents, in: *Proceedings of the IEEE Military Communication Conference, MILCOM 2003*, vol. 1, Boston, MA (October 13–16, 2003), pp. 705–710.
46. P. Venkitasubramaniam, S. Adireddy and L. Tong, Sensor networks with mobile agents: Optimal random access and coding, IEEE Journal on Selected Areas in Communications 22(6) (2004) 1058–1068.
47. P. Venkitasubramaniam, Q. Zhao and L. Tong, Sensor networks with multiple mobile access points, in: *Proceedings of the 38th Annual Conference on Information Sciences and Systems, CISS 2004*, Princeton, NJ (March 17–19, 2002).
48. G. Wang, G. Cao and T.F. La Porta, A bidding protocol for deploying mobile sensors, in: *Proceedings of the 11th IEEE International Conference on Networks Protocols, ICNP 2003*, Atlanta, GA (November 4–7, 2003), pp. 315–324.
49. G. Wang, G. Cao and T.F. La Porta, Movement-assisted sensor deployment, in: *Proceedings of IEEE INFOCOM 2004*, vol. 4, Hong Kong (March 7–11, 2004), pp. 2469–2479.
50. G. Wang, G. Cao and T.F. La Porta, Proxy-based sensor deployment for mobile sensor networks, in: *Proceedings of the First IEEE International Conference on Mobile Ad Hoc and Sensor Systems, MASS 2004*, Fort Lauderdale, FL (October 25–27, 2004), pp. 493–502.
51. G. Wang, G. Cao, T.F. La Porta and W. Zhang, Sensor relocation in mobile sensor networks, in: *Proceedings of IEEE INFOCOM 2005*, vol. 4, Miami, FL (March 13–17, 2005), pp. 2302–2312.
52. W. Wang, V. Srinivasan and K.-C. Chua, Using mobile relays to prolong the lifetime of wireless sensor networks, in: *Proceedings of the 11th Annual ACM/SIGMOBILE International Conference on Mobile Computing and Networking, MobiCom 2005*, Cologne Germany (August 28–September 2, 2005), pp. 270–283.
53. Z.M. Wang, S. Basagni, E. Melachrinoudis and C. Petrioli, Exploiting sink mobility for maximizing sensor networks lifetime, in: *Proceedings of the 38th Hawaii International Conference on System Sciences*, Big Island Hawaii (January 3–6, 2005).
54. H.L. Xuan and S. Lee, A coordination-based data dissemination protocol for wireless sensor networks, in: *Proceedings of the 2004 Intelligent Sensors, Sensor Networks and Information Processing Conference, ISSNIP 2004*, Melbourne Australia (December 14–17, 2004), pp. 13–18.
55. F. Ye, H. Luo, J. Cheng, S. Lu and L. Zhang, A two-tier data dissemination model for large scale wireless sensor networks, in: *Proceedings of the 8th ACM Annual International Conference on Mobile Computing and Networking, MobiCom 2002*, Atlanta, GA (September 23–28, 2002), pp. 148–159.
56. W. Zhao, M.H. Ammar and E.W. Zegura, A message ferrying approach for data delivery in sparse mobile ad hoc networks, in: *Proceedings of the 5th ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc 2004*, Roppongi Hills, Tokyo, Japan (May 24–26, 2004), pp. 187–198.
57. W. Zhao, M.H. Ammar and E.W. Zegura, Controlling the mobility of multiple data transport ferries in a delay-tolerant network, in: *Proceedings of IEEE Infocom 2005*, vol. 2, Miami, FL (March 13–17, 2005), pp. 1407–1418.
58. M. Zorzi and R.R. Rao, Geographic random forwarding (GeRaF) for ad hoc and sensornetworks: Energy and latency performance, IEEE Transactions on Mobile Computing 2(4) (2003) 349–365.
59. M. Zorzi and R.R. Rao, Geographic random forwarding (GeRaF) for ad hoc and sensor networks: Multihop performance, IEEE Transactions on Mobile Computing 2(4) (2003) 337–348.

Stefano Basagni holds a Ph.D. in electrical engineering from the University of Texas at Dallas (December 2001) and a Ph.D. in computer science from the University of Milano, Italy (May 1998). He received his B.Sc. degree in computer science from the University of Pisa, Italy, in 1991. Since Winter 2002 he is on faculty at the Department of Electrical and Computer Engineering at Northeastern University, in Boston, MA. From August 2000 to January 2002 he was professor of computer science at the Department of Computer Science of the Erik Jonsson

School of Engineering and Computer Science, The University of Texas at Dallas.

Dr. Basagni's current research interests concern research and implementation aspects of mobile networks and wireless communications systems, Bluetooth and sensor networking, definition and performance evaluation of network protocols and theoretical and practical aspects of distributed algorithms.

Dr. Basagni has published over four dozens of referred technical papers and book chapters. He is also co-editor of two books. Dr. Basagni served as a guest editor of the special issue of the Journal on Special Topics in Mobile Networking and Applications (MONET) on Multipoint Communication in Wireless Mobile Networks, of the special issue on mobile ad hoc networks of the Wiley's Interscience's Wireless Communications & Mobile Networks journal, and of the Elsevier's journal Algorithmica on algorithmic aspects of mobile computing and communications.

Dr. Basagni serves as a member of the editorial board and of the technical program committee of ACM and IEEE journals and international conferences. He is a senior member of the ACM (including the ACM SIGMOBILE), senior member of the IEEE (Computer and Communication societies), and member of ASEE (American Society for Engineering Education).

Alessio Carosi received the M.S. degree "summa cum laude" in Computer Science in 2004 from Rome University "La Sapienza." He is currently a Ph.D. candidate in Computer Science at Rome University "La Sapienza." His research interests include protocols for ad hoc and sensor networks, underwater systems and delay tolerant networking.

Emanuel Melachrinoudis received the Ph.D. degree in industrial engineering and operations research from the University of Massachusetts, Amherst, MA. He is currently the Director of Industrial Engineering and Associate Chairman of the Department of Mechanical and Industrial Engineering at Northeastern University, Boston, MA. His research interests are in the areas of network optimization and multiple criteria optimization with applications to telecommunication networks, distribution networks, location and routing. He is a member of the Editorial Board of the International Journal of Operational Research. He has published in journals such as Management Science, Transportation Science, Networks, European Journal of Operational Research, Naval Research Logistics and IIE Transactions.

Chiara Petrioli received the Laurea degree "summa cum laude" in computer science in 1993, and the Ph.D. degree in computer engineering in 1998, both from Rome University "La Sapienza," Italy. She is currently Associate Professor with the Computer Science Department at Rome University "La Sapienza." Her current work focuses on ad hoc and sensor networks, Delay Tolerant Networks, Personal Area Networks, Energy-conserving protocols, QoS in IP networks and Content Delivery Networks where she contributed around sixty papers published in prominent international journals and conferences. Prior to Rome University she was research associate at Politecnico di Milano and was working with the Italian Space agency (ASI) and Alenia Spazio. Dr. Petrioli was guest editor of the special issue on "Energy-conserving protocols in wireless Networks" of the ACM/Kluwer Journal on Special Topics in Mobile Networking and Applications (ACM MONET) and is associate editor of IEEE Transactions on Vehicular Technology, the ACM/Kluwer Wireless Networks journal, the Wiley InterScience Wireless Communications & Mobile Computing journal and the Elsevier Ad Hoc Networks journal. She has served in the organizing committee and technical program committee of several leading conferences in the area of networking and mobile computing including ACM Mobicom, ACM Mobihoc, IEEE ICC,IEEE Globecom. She is member of the steering committee of ACM Sensys and of the international conference on Mobile and Ubiquitous Systems: Networking and Services (Mobiquitous) and serves as member of the ACM SIGMOBILE executive committee. Dr. Petrioli was a Fulbright scholar. She is a senior member of IEEE and a member of ACM.

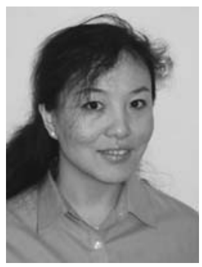

Z. Maria Wang received her Bachelor degree in Electrical Engineering with the highest honor from Beijing Institute of Light Industry in China, her M.S. degree in Industrial Engineering/Operations Research from Dalhousie University, Canada and her Ph.D. in Industrial Engineering/Operations Research from Northeastern University, Boston. She served as a R&D Analyst for General Dynamics. Currently MS. Wang serves as an Optimization Analyst with Nomis Solutions, Inc.