

Pro-active Strategies for the Frugal Feeding Problem in Wireless Sensor Networks

Elio Velazquez, Nicola Santoro, and Mark Lanthier

School of Computer Science, Carleton University, Canada
evelazqu, santoro, lanthier@scs.carleton.ca

Abstract. This paper proposes a pro-active solution to the Frugal Feeding Problem (FFP) in Wireless Sensor Networks. The FFP attempts to find energy-efficient routes for a mobile service entity to rendezvous with each member of a team of mobile robots. Although the complexity of the FFP is similar to the Traveling Salesman Problem (TSP), we propose an efficient solution, completely distributed and localized for the case of a fixed rendezvous location (i.e., service facility with limited number of docking ports) and mobile capable entities (sensors). Our pro-active solution reduces the FFP to finding energy-efficient routes in a dynamic Compass Directed unit Graph (CDG). The proposed CDG incorporates ideas from forward progress routing and the directionality of compass routing in an energy-aware unit sub-graph. Navigating the CDG guarantees that each sensor will reach the rendezvous location in a finite number of steps. The ultimate goal of our solution is to achieve energy equilibrium (i.e., no further sensor losses due to energy starvation) by optimizing the use of the shared resource (recharge station). We also examine the impact of critical parameters such as transmission range, cost of mobility and sensor knowledge in the overall performance.

1 Introduction

The problem of achieving continuous operation in a robotic environment by refueling or recharging mobile robots has been the focus of attention in recent research papers. In particular, [12, 13] present this problem as the Frugal Feeding Problem (FFP), for its analogy with occurrences in the animal kingdom. The FFP attempts to find energy-efficient routes for a mobile service entity, also called “tanker”, to rendezvous with every member of a team of mobile robots. The FFP has several variants depending on where the “feeding” or refueling of the robots takes place: at each robot’s location, at a predefined location (e.g., at the tanker’s location) or anywhere. Regardless of which variant is chosen, the problem is to ensure that the robots reach the rendezvous location without “dying” of energy starvation during the process.

In a Wireless Sensor Network (WSN) deployment, the sensors will eventually deplete their batteries and loss of coverage would occur. Some approaches to cope with an eventual loss of coverage attempt to extract energy from the environment to extend network lifetime [19, 20]. Others explore the use of mobile entities (e.g.,

robots, actuators, service stations) in conjunction with clustering techniques [16, 24, 11, 22]. In general, energy management strategies can be categorized into two groups: *cluster-based* approaches (e.g., [18, 14, 28, 9]) or *mobility-based* approaches (e.g., [17, 27, 15, 26, 11]) with some degree of overlap.

In this paper we study the FFP in a wireless sensor network scenario where mobility capabilities are added to the sensors and static recharge facilities are deployed throughout the sensing area. In this variant of the FFP, the responsibility for maintaining the overall health of the network is shifted to the sensor side, whereas the service facilities play a more passive role. The rendezvous between sensors and facilities should take place at the closest facility’s original position, which is static. The maximum number of sensors that can rendezvous with a facility at any given time is determined by the number of docking ports or recharge sockets available at the facility.

According to the FFP terminology introduced in [12], our problem can be seen as the “tanker absorbed” version of the FFP. The rendezvous between the service facility (i.e., tanker robot) and the mobile robots (i.e., mobile sensors in our case) takes place at the current location of the service facility. The location of the service facility is known a priori and the problem is reduced to finding energy efficient routes to reach the facility. Another characteristic of our scenario is that the sensors are static in nature. That is, they have been deployed and have been assigned specific tasks. Therefore, their movement to the rendezvous location will create coverage holes that should be kept minimal.

The sensors need to communicate and coordinate their actions in order to achieve a common goal (i.e., continuous sensing operation without losses due to energy starvation). Furthermore, sensors should coordinate their moves in a loop-free manner so the intended destination (i.e., recharge station) is reached in a finite number of moves or steps. The ultimate goal of the FFP is to reach a state of energy equilibrium where there are no further sensor losses. This work also examines some underlying topologies that guarantee a loop free mobility strategy as well as the network parameters needed to achieve the state of equilibrium.

1.1 Related Work

In the FFP, as described in [12], specialized robots (called tankers) have to rendezvous with mobile robots to refuel or recharge them. The main goal is to minimize the amount of fuel (energy) required to move the robots and tankers to the rendezvous locations. The problem can have several variants: 1) the rendezvous can take place at the robot’s location. The robots in need of energy do not move but instead wait for the refueling tanker to come to their rescue. This is called: the robot-absorbed case. 2) the rendezvous takes place at the tanker’s location and the robots should move to the tanker’s original location. This is called tanker-absorbed case. 3) the rendezvous takes place at locations that do not coincide with the initial robot or tanker locations. The FFP also has a combinatorial component pertaining to the order in which the robots should be recharged. Finding a solution to the FFP that guarantees that no robots die of energy starvation is an NP-Hard problem (as shown in [12]).

The problem of determining where to place a docking station or recharger is examined by [4]. In this case, a team of mobile robots have the specific task of transporting certain items from a pick-up location to a drop-off location. To be able to work for a prolonged period of time, the robots should interrupt their work and visit the recharge station periodically (i.e., tanker-absorbed FFP). Their solution is to place the charger station close enough to the path followed by the robots but without causing interference to the robot's movements.

Examples of the robot-absorbed FFP can be found in [3, 5, 21]. In both cases, a charger robot is responsible for delivering energy to a swarm of robots. The recharging strategy is completely reactive (i.e., robots are only recharged when they become out of service and cannot move). In the scenario described in [3], the charger robot is equipped with several docking ports. However, the charger robot can travel to recharge a needed robot only if none of the docking ports are occupied, assuming that several depleted robots need to be close by in order to be recharged simultaneously. The simulation results presented in [5] showed that in a network with 64 robots and one charger station with only one docking port; there will be a large number of robots either abandoned or dead due to battery depletion. However, increasing the number of docking ports to 2, affects the performance dramatically by decreasing the number of robot deaths and improving the exploring/dead time ratios. The solution presented in [21] creates clusters based on the number of available chargers. The experimental results with this approach show that a network with 76 sensors deployed in an area of $1000 \times 1000 m^2$ requires at least 3 chargers to keep the network alive. The network is considered dead when more than 50% of the sensors die due to battery depletion.

All the aforementioned scenarios satisfy the necessary conditions for mobile robots to be able to recharge themselves as presented in [3]. For example, the robots should be able to monitor their energy levels and detect when it is time to recharge. Second, they should be able to locate and move towards a charging station. Finally, there should be a mechanism for energy transfer either by docking or plugging in to the charging station or via wireless recharging at short distances (e.g. [1, 2, 17]).

1.2 Contributions

This paper proposes a pro-active solution to the Frugal Feeding Problem (FFP) in Wireless Sensor Networks. We propose an efficient solution, completely distributed and localized for the case of a fixed rendezvous location (i.e., service facility with limited number of docking ports) and mobile sensors. In particular, we propose to reduce the tanker-absorbed FFP with a fixed rendezvous location in a sensor network of arbitrary topology to finding energy-efficient routes in a dynamic Compass Directed unit Graph (CDG). We prove that energy-aware mobility strategies using the CDG are loop-free, guaranteeing that the sensors will reach the recharge station within a finite number of moves. The experimental analysis of our solution confirms that energy equilibrium (i.e., no further losses due to energy starvation) can be achieved in a network of 100:1 sensor/station

ratio with one station containing two docking ports. Our experiments also examine the impact of critical parameters such as transmission range, cost of mobility and station role.

The main differences between our proposed solution to the FFP and the existing literature in the area of autonomous robot recharging are: 1) Our solution is completed distributed and localized; there is no need for an entity with global knowledge. Sensors are only aware of their immediate neighbors and the location of the closest facility. 2) Our approach is completely pro-active. The sensors act before their batteries reach a critical level to minimize coverage holes by making the shortest possible trip to the recharge station. 3) The algorithms for route selection and logical topologies used are dynamic and adaptive.

2 Pro-active Solution to the Facility-Absorbed FFP

Our pro-active solution to the FFP in the sensor network scenario is built from two main components: mobile capable sensors and static recharging facilities. The general requirement for our theoretical model is to maximize the network operating life by the autonomous recharging of low energy sensors. However, the ultimate goal is to achieve a state of equilibrium where no further losses are reported and to accomplish this with the minimum amount of resources. In general, the model includes the following key components: 1) A set of N sensors, $S = \{s_1, \dots, s_N\}$ randomly distributed in an area of unspecified shape. 2) A randomly located static recharge facility F (i.e., rendezvous location). The facility is equipped with a fixed number of recharging ports or sockets. This represents the maximum number of simultaneous sensors at the rendezvous location.

It is assumed that sensors can determine their own positions by using GPS or some other localization method. Sensors can communicate with other sensors within their transmission range R and they all move at the same speed. The distance to the closest facility should be within the sensors' mobility range to guarantee a successful round-trip to the station with one battery charge. All communications are asynchronous; there is no global clock or centralized entity to coordinate communications or actions.

We consider the sensors to be static in terms of their sensing requirements. In other words, from the point of view of the application (i.e., functional requirements), the sensors are static and placed at a specific set of coordinates. However, they all have the capability of moving if they decide to go to the service station to recharge their batteries. Consequently, a pro-active behavior implies that the sensors decide to act before their batteries reach a critical level. The general idea is that sensors will try to get closer to the rendezvous location by swapping positions with other sensors that are closer to the station and eventually make the shortest possible trip when their batteries reach a critical level.

Every time a sensor visits the recharge station, a coverage hole is created. The duration of the hole depends on the recharging time plus the length of the round-trip. In order to minimize coverage holes sensors will attempt a gradual

approach towards the rendezvous location by swapping positions with higher energy sensors. The operating life of a sensor is divided in three stages depending on its battery status: 1) BATTERY_OK or normal operation, 2) BATTERY_LOW or energy-aware operation and 3) BATTERY_CRITICAL or recharge-required operation. A sensor in a BATTERY_OK state will perform its regular sensing functions as well as accept any swapping proposal from other sensors with less energy. When the battery level falls below a fixed threshold, the sensor switches its state to a more active BATTERY_LOW state. In this state, the sensor will start its migration towards the service station, proposing swapping operations to sensors with higher energy levels. Finally, a sensor in the BATTERY_CRITICAL state will contact the service station and wait until a socket or docking port has been secured, then it will travel to the station and recharge (see Figure 1).

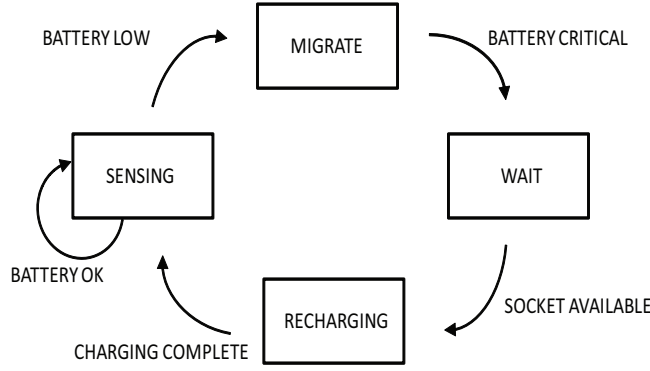


Fig. 1. A sensor's life cycle.

In this life cycle, it is the migration behavior that is of interest. The objective of the sensor during migration is to reach the recharge facility in an effective timely manner, while relying solely on local information. This can be done by allowing the sensor to explore energy-aware routes leading to the recharge facility. The chosen routes are based on a logical Compass Directed unit Graph (CDG).

Definition 1. A graph $G = (V, E)$ with vertices $V = \{v_1, \dots, v_N\}$ and edges $E = \{(v_i, v_j)\}$ with $1 \leq i < j \leq N$ is called a *Unit Disk Graph (or Unit Graph)* if $d(v_i, v_j) \leq R$ where d is the Euclidean distance between the sensors and R is the transmission range.

Definition 2. A graph $G' = (V' \cup F, E')$ with $V' \subseteq V$ and $E' \subseteq E$ is called *Compass Directed unit Graph (CDG)* if \forall pair of sensors $S_i, S_j \in V'$ and recharge facility F , the following conditions are satisfied:

$$\text{Unit graph criterion: } d(S_i, S_j) \leq R \quad (1)$$

where d denotes the Euclidean distance and R is the transmission range.

$$\text{Proximity criterion: } d(S_j, F) < d(S_i, F) \text{ and } d(S_i, S_j) < d(S_i, F) \quad (2)$$

$$\text{Directionality criterion: } \forall S_i, S_j \text{ pair, } \exists S_{jp} \text{ such that } S_j \vec{S}_{jp} \cdot S_i \vec{F} = 0 \text{ and } d(S_i, S_{jp}) + d(S_{jp}, F) = d(S_i, F) \quad (3)$$




Fig. 2. Compass Directed unit Graph

Routing algorithms use the hop count as the metric to measure effectiveness. In our case, the hop count would be equivalent to the number of swapping operations between sensors in our CDG. Our solution to the FFP can be divided into two main stages: 1) the construction of the CDG and 2) the incremental swapping approach (i.e., migration) towards the rendezvous location.

2.1 Creating the CDG

Figure 2 shows an example of the proposed CDG for three sensors A,B,C and a facility F. In the first stage of the algorithm, it is assumed that all sensors have the required levels of energy to construct the CDG. The process is rather simple and can be summarized by the following actions:

1. Sensors position themselves at some initial fixed location that depends on the task at hand.
2. Sensor A sends a NEIGHBOR_REQUEST broadcast message inviting other sensors to participate.
3. Upon receiving a NEIGHBOR_REQUEST message from sensor A, immediate neighbors verify the neighboring criteria according to the following rules:
 - a) Proximity: $d(A, F) > d(B, F)$ and $d(A, B) < d(A, F)$.
 - b) Directionality: For example, B and C are neighbors of A if the corresponding projections B_p and C_p on line \overline{AF} intersect the line segment \overline{AF} .
4. If the conditions a) and b) are met, then sensors B and C send a NEIGHBOR_ACCEPT message. Otherwise they send a NEIGHBOR_DENY message.

In order to save energy, sensor A will then try to deviate as little as possible from the direction of the recharge station F . That is, sensor A will try to minimize the angle $\angle BAB_p$. Therefore, all the sensors that satisfy the conditions a) and b) are ranked according to the following function: $f(S_i, S_j) = \left\{ d(S_i, S_j) + \frac{d(S_i, S_{jp})}{d(S_i, S_j)} \right\}$ where S_i, S_j are the neighboring sensors, d is the Euclidean distance, F is the recharge station and S_{jp} is the projection of S_j on the line segment $\overline{S_i F}$.

At the end of this phase each sensor will have two routing tables: one containing its children (i.e., sensors from which `NEIGHBOR_ACCEPT` messages were received) with their corresponding ranking and a second table containing its parents (i.e., sensors to which `NEIGHBOR_ACCEPT` messages were sent). The routing tables are just partial maps of the network indicating the position of the children and parents.

2.2 Migration Strategy

The second stage of the algorithm starts when sensors change their state from `BATTERY_OK` to `BATTERY_LOW` as a result of their battery levels falling below the first threshold. Once a sensor enters this state, it will try to get closer to the facility by making a series of one-hop swaps with its graph neighbors. Some of the most relevant sensor interactions in this stage can be summarized by Algorithm 1.

In an ideal system, all sensors will reach the `BATTERY_CRITICAL` state when they are exactly at one-hop distance from the rendezvous location. When the trip to the recharge station is made from a one-hop position (i.e., there are no graph neighbors), we call this a “one-hop run” or “optimal run”. Contrarily, if the final trip is made from any other location, it is called a “panic run”.

2.3 Properties of the CDG

There are two important properties of the CDG (i.e., dynamic and self-correcting) that can be explained by the following scenarios. Both scenarios may cause situations where the information in the neighboring tables is obsolete.

- Scenario 1: Simultaneous swapping. As part of the swapping process, the participating sensors exchange their neighboring information, that is, their corresponding children and parent tables. However, since multiple swapping operations may occur at the same time, when a sensor finally arrives at the position occupied by its swapping partner, the information in its neighboring tables may be out-of-date.
- Scenario 2: Sensor recharging. While this process takes place, other sensors may be swapping positions. Once the recharging process is finished, the sensor returns to its last known position. However, the structure of the network around it has changed. This situation is even more evident when trips to the facility are made from distances of more than one hop as a result of “panic runs”.

The solution to these problems is to define the neighboring information as position-based tables, where the important factor is the relative position of the neighbors and not their corresponding IDs. The information of the actual sensors occupying the positions is dynamic. In other words, a sensor knows that at any given point in time it has n children at positions $(x_1, y_1) \dots (x_n, y_n)$ and p parents at positions $(x'_1, y'_1) \dots (x'_p, y'_p)$. This information is static and will not be modified. However, the identity of the sensors occupying the positions is dynamic and

Algorithm 1 Excerpt of the migration algorithm for sensor S to facility F

```

(* In State BATTERY_OK : *)
if BATTERY_LEVEL < BATTERY_LOW_THRESHOLD then
  rank=1
  become BATTERY_LOW
end if
(* In State BATTERY_LOW : *)
if BATTERY_LEVEL < BATTERY_CRITICAL_THRESHOLD then
  send RECHARGE_REQUEST message to Facility
  become BATTERY_CRITICAL
else
  while rank ≤ numberOfNeighbors do
    send SWAP_REQUEST to sensor with rank: rank
    become WAIT_FOR_SWAP_REPLY
  end while
end if
(* In State WAIT_FOR_SWAP_REPLY_STATE : *)
if receiving SWAP_ACCEPT from  $S_i$  then
  move to  $S_i$ 
  send SWAP_COMPLETE
  rank=1
  become BATTERY_LOW
end if
if receiving SWAP_DENY from  $S_i$  then
  rank = rank + 1
  become BATTERY_LOW
end if
(* In State BATTERY_CRITICAL : *)
if receiving RECHARGE_ACCEPT then
  lastPosition = currentPosition
  move to Facility
  recharge
  move to lastPosition
  send SENSOR_RECHARGED message
  become BATTERY_OK
end if

```

will get updated every time a swapping operation occurs. The mechanism to detect changes in the routing tables is triggered by sending a *SWAP_COMPLETE* message. When two neighboring sensors successfully complete a swapping operation, they will announce their new positions by sending *SWAP_COMPLETE* messages. Sensors within the transmission range that listen to this message will verify whether any of the positions involved in the exchange belongs to their routing tables and update the appropriate entry with the new occupant of that position.

On the other hand, a sensor returning from the service station (e.g., scenario 2) needs to re-discover the new occupants of its routing tables. This process is

initiated by a `SENSOR_RECHARGED` message sent by the newly recharged sensor as soon it reaches its last known position on the network. Potential children and parents, upon receiving this message, will reply with `CHILD_UPDATE` and `PARENT_UPDATE` messages accordingly. This process is also used for parents to update their information about the energy levels of this newly recharged sensor.

These two important properties, along with a neighboring criteria that incorporates ideas from forward progress and compass routing [23, 10, 8] in an energy-aware unit graph, ensure the following lemma:

Lemma 1. *The swapping-based pro-active solution to the FFP guarantees that all sensors reach the rendezvous location within a finite number of swapping operations.*

Proof. Let $G = (V, E)$ be a CDG with a set of vertices $V = \{S_1, \dots, S_N, F\}$ where S_i , $1 \leq i \leq N$ represent sensors and F denotes the rendezvous location. Let E be a set of edges of the form $S_i \rightarrow S_j$ where S_j is neighbor of S_i . By definition, G satisfies the conditions of proximity (2) and directionality (3).

Without loss of generality, we can assume that for any path $P_i = \langle S_i, \dots, S_K, F \rangle$ leading to the recharge station F , with $1 \leq i < K \leq N$, the sub-path containing the sensors $\langle S_i, \dots, S_K \rangle$ does not contain any cycles. This claim can be proved by contradiction.

Let us assume that the rendezvous location cannot be reached. This means that at some point during the execution of the algorithm a given sensor finds itself in a loop (i.e., a cycle C of arbitrary length L is found). Let $C = \{S_i, S_{(i+1)}, \dots, S_{(L-1)}\} \cup \{S_L, S_i\}$ with $1 \leq i < L \leq N$. If such a cycle C exists, sensor S_i must be neighbor of sensor S_L which means that $d(S_i, T) < d(S_L, T)$. This contradicts the proximity criterion (2)(triangular inequality). Hence, the Lemma holds. \square

During the algorithm, the facility plays a rather passive role. The facility's responsibilities are limited to keeping a queue of waiting sensors ranked by their energy levels and notify the sensors when a socket or docking port becomes available. In a passive scenario, a socket becomes available when the sensor has reached 100% of its battery level and sends a `RECHARGE_DONE` message to the facility. In an active scenario, the facility does not have to wait for the sensor's battery to be 100% recharged. In this case, a sensor will notify the facility when its battery has reached an operational level (e.g., 75%). Consequently, the facility could halt the charging process by sending a `TERMINATE_RECHARGE`, if there are other sensors waiting in line.

3 Experimental Results

Previous work on energy consumption of wireless sensor networks and protocols such as 802.11, have found that the energy required to initiate communication

is not negligible. In particular, loss of energy due to retransmissions, collisions and acknowledgments is significant [6, 7]. Therefore, protocols that rely on periodic probe messages and acknowledgments are considered high cost. For these reasons, the design of our algorithm and related coordination had to be flexible enough to avoid the use of probe messages and complicated state-full protocols. It is also noted in the literature that energy consumption of sensors in idle state can be as large as the energy used when receiving data [7]. On the other hand, the energy used in transmitting data is between 30-50% more than the energy needed to receive a packet. These differences, between the energy needed to perform the basic operations and the percentage of battery usage, vary depending on the communication protocols, hardware and type of battery used.

A common problem faced by any solution involving mobile entities is that of determining a way to accurately represent the cost of energy spent when moving from one location to another. Locomotion cost depends on many factors such as weight of the electronic components, irregularities in the terrain, obstacles, etc. For simplicity, in [12], the weighted Euclidian distance between the origin and destination is used as the cost of relocating a robot. In this paper we consider an experimental setting based on real robots deployed in a controlled environment. The goal of the experiments was to identify the impact of basic operations (i.e., communication, sensing, locomotion, idle operation) on the overall battery life.

The experiments were based on the PropBot 2.0 mobile robot (Figure 3(a)) developed in the School of Computer Science Robotics Lab at Carleton University. The robot's hardware specification includes: Parallax Propeller Microprocessor (with 8 processors), Parallax Continuous Rotation Servos, CUMCam camera, three Sharp GP2Y0D810Z0F Digital Distance Sensors and Nubotics WW-01 Encoders. Communications use the Parallax EasyBluetooth module and batteries are custom-made 6v battery packs using 2600mAh NIMH AA cells (Figure 3(b)). A single mobile robot was deployed in an area of 2m x 1.5m and tests were performed to determine battery drain under the following conditions: 1) idle state, 2) continuous movement, 3) communication 4) sensor usage.




Fig. 3. Experimental equipment

Our preliminary results show that energy spent in communications (i.e., send/receive) is 25% more than the battery drain in the idle state. Battery drain under perpetual movement (i.e., locomotion costs) is almost twice as much as communication cost. Sensing with the CUMCam was among the most costly operations and the battery recharge was 14x faster than battery drain in the idle state. These findings were incorporated into our simulation scenarios to study the impact of critical variables on the overall solution.

The simulation scenarios are implemented in Omnet++ along with the mobility framework extension [25]. For all experiments, the sensors and charging facilities were randomly placed in an area of $1000 \times 1000 m^2$. The analysis of our simulated results centers on three important aspects of the solutions:

- 1) Whether or not a state of equilibrium is achieved and the number of sensor losses until such condition is met.
- 2) Impact of several variables such as: transmission range, mobility cost, etc.
- 3) Role played by the charging station: passive vs. active.

In all cases the quality of the strategy is measured in terms of optimal runs vs. panic runs. Constant cost values are assigned to each basic operation (i.e., send, receive, idle and move). Initial values for these operations are based on the observations with the PropBot robot as well as some of the experiences found in the literature [6, 7].

3.1 Sensor Losses Over Time

Our first test attempted to determine whether our pro-active solution to FFP reaches a state of equilibrium and to measure the cumulative number of losses until this condition is met. In other words, it measured the number of sensor losses over time until the system reached a state where no more losses were reported. Figure 4(a) shows the result of a simulation involving 100 sensors and one service facility. The facility is equipped with two sockets which allow only two sensors to be recharged at the same time. A series of 30 experiments with different random deployments were run for 10^6 simulation seconds. The sensor transmission range is fixed at 100m and the energy ratio for sending/receiving a packet is set to a constant ($E : E/2$). Locomotion costs were based on the weighted Euclidean distance with a weight factor of $1/5E$ per meter traveled. Confirming our expectations, our algorithm reached the state of equilibrium for all the random deployments. In comparison, the work of [5] and [21], required two and 3 stations or actors, respectively, to maintain a live network (50% or more sensors remain after equilibrium was reached). In our case, equilibrium was achieved with 1 facility with two docking ports for a similar network size and over 80% of network survivability.

3.2 Transmission Range and Mobility Cost

The second experiment was designed to verify the impact of the sensor's transmission range on the overall performance. The characteristics of the network

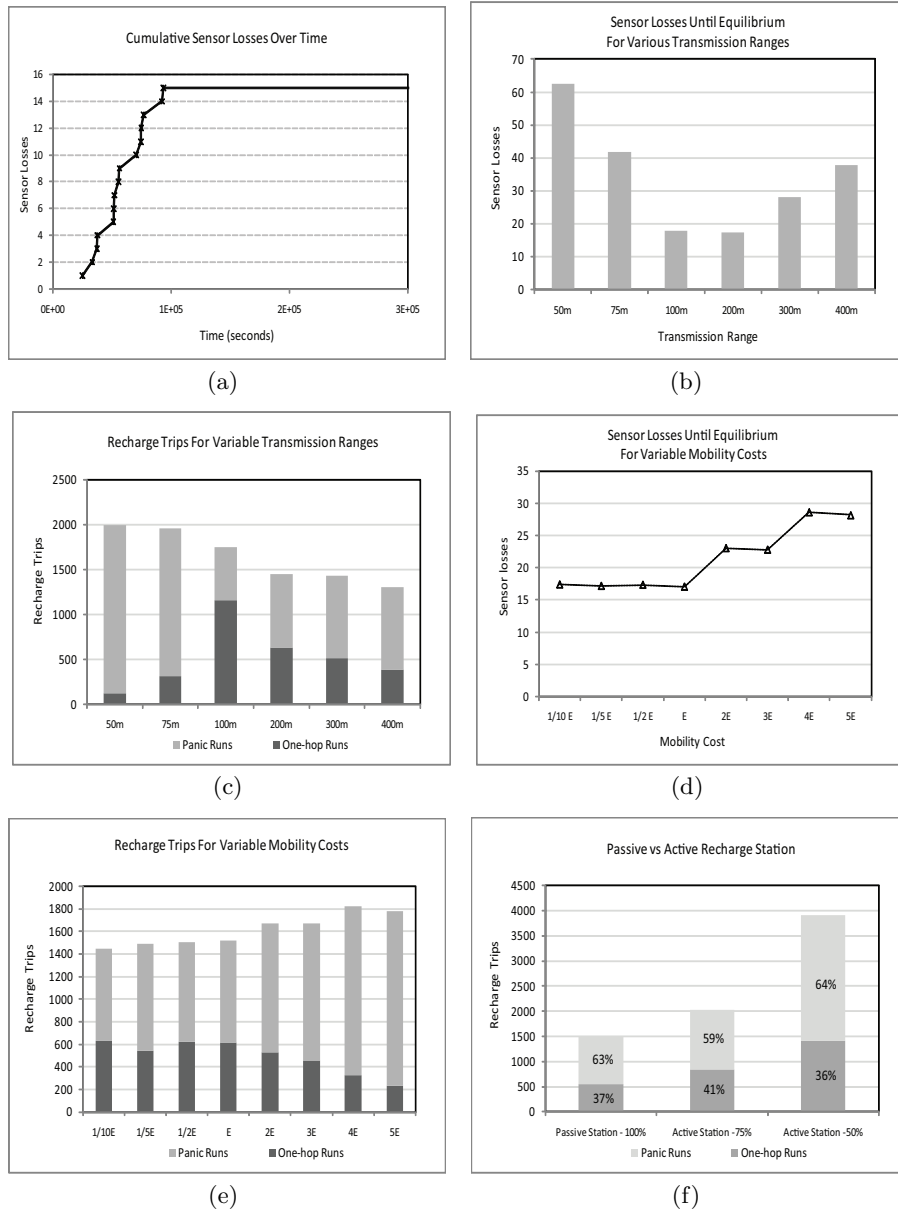


Fig. 4. Experimental results

were the same as in the previous test and the experiments are run the same length of time. The only difference is that the transmission range was varied from 50m, 75m, 100m, 200m, 300m and 400m. Figure 4(b) shows the cumulative number of sensor losses until equilibrium for each range value. In a deployment

of $1000 \times 1000 m^2$ a transmission range of 50m was too restrictive, which means that most of the sensors were isolated and the number of immediate neighbors in the CDG was too small to guarantee a gradual approach towards the recharge location. Another interesting observation is that by increasing the transmission range, the number of losses decreased dramatically. However, for larger ranges (300m and 400m) there was a decline on the overall performance, since many neighbors were discovered resulting in an added overhead to maintain more information per sensor as well as additional interactions due to update messages as result of successful swapping and recharging operations.

Figure 4(c) shows the quality of the solution in terms of one-hop runs vs. panic runs. In an ideal system, our solution should reach the state of equilibrium using one-hop runs only. As expected, for a transmission range of 50m, most of the trips could be considered panic runs since there is almost no migration due to the lack of one-hop neighbors. The best breakdown between one-hop and panic runs occurs with 100m range. However, there are more visits to the recharge location, when compared to the 200m, 300m and 400m cases. Although there is no clear explanation to this phenomenon, one can argue that there is a trade-off between the total number of recharge trips and the breakdown between one-hop vs. panic runs. In a panic run situation, a sensor travels from a more distant location and after having been recharged, it needs to travel further in order to return to its initial location. This situation creates a coverage hole that would last for a longer period of time, when compared to a one-hop run. However, more one-hop recharge trips also means more coverage holes but for shorter periods of time.

The next experiment explored the impact of locomotion cost on the number of losses until equilibrium was reached as well as the distribution and number of recharge trips. The network setup remained the same with the transmission range fixed at 100m. Figure 4(d) shows the number of losses until equilibrium for several mobility costs. The cost function is based on the weighted distance traveled by the sensors, with the weight constant w defined as a function of the energy spent to send a packet. For example: if E is the energy spent to send a packet over the 100m range, then for each meter traveled, the sensor will spend wE units of energy, where $w \in \{1/10, 1/5, 1/2, 1, 2, 3, 5\}$. In other words, the energy spent to move the robot 100m, ranges from 10x to 500x the energy required to send a packet over the same distance. In particular, the values observed for the ProbBot robot fluctuated around 54x the communication energy.

The simulation results show that as the locomotion costs increase (in relation to the transmission cost) so does the number of sensor losses until equilibrium. The trend seems to be closer to a step function with clear discrete increments at some values. Another observation is that despite the increase in the number of sensor losses, the network survivability is still over 70%, even for the worst case.

In terms of the quality of the solution in the variable mobility cost scenario, Figure 4(e) shows the same step function behavior for the total number visits to the station. However, there is a significant degradation on the number of one-

hop trips as the locomotion cost increases because a larger number of sensors fall into the BATTERY_CRITICAL state before completing their migration.

3.3 Passive vs Active Charging Station

The last experiment examined the case where the recharge station was given a more active role in order to minimize sensor losses while waiting for an available socket. Figure 4(f) shows the comparison between passive vs. active charging stations while using the same sensor deployment as the previous tests. These particular tests used a fixed transmission range of 100m, a mobility cost factor of $1/5E$ per meter traveled and a recharge rate 10x faster than battery drain in idle state. In this experiment, the sensors notified the station when their batteries reached 100%, 75% and 50% of charge. By giving a more active role to the facility, the network survivability at the state of equilibrium was improved from 80% to close to 90% for the 75% recharge case. The number of one-hop trips was also improved from 37% to 41% for the 75% sensor recharge case. However, the number of recharge trips increased by 30%. Recharging the batteries to 50% of their capacity almost doubled the number of recharge visits when compared to the 75% case. The number of one-hop trips also decreased but the network survivability remained close to 90%. Once again there is a trade-off between creating temporary coverage holes produced by additional recharge trips and permanent coverage holes produced by sensor losses due to battery depletion.

4 Conclusions and Future Work

In this work we have presented an efficient, completely distributed and localized solution for the facility absorbed Frugal Feeding Problem (FFP). Our novel solution recommends taking a pro-active approach to energy restoration based on the construction of a Compass Directed Graph (CDG) and a swapping-based incremental approach towards the rendezvous location. In summary, our pro-active solution has the following properties:

- 1) The proposed CDG guarantees that sensors will reach the rendezvous location within a finite number of swapping operations. The trajectory is loop-free.
- 2) All decisions made by the sensors regarding the next swapping operation are based on local knowledge (i.e., the algorithms are completely distributed and localized).
- 3) The proposed CDG and the incremental swapping algorithm are dynamic and self-correcting: neighboring information is updated any time a successful swapping or recharge operation takes place.

The experimental analysis of our novel pro-active solution to the FFP shows that for networks of 100:1 sensor/facility ratio, a state of energy equilibrium can be reached with over 80% network survivability. The simulations also expose several trade-offs between key variables (i.e., transmission range, locomotion cost) and the quality of the overall solution in terms of optimal and panic visits

to the facility. The simulations also show that when giving the facility a more active role during the recharging process, the state of equilibrium can be reached with close to 90% network survivability. The breakdown between optimal and panic runs is also improved but there is an increase in the overall number of recharge trips.

Future enhancements to this work may explore in more detail the impact of mobility. Locomotion costs are very dependent on physical conditions, hardware specifications, battery technology, etc. This may involve the use of the PropBot mobile robots in larger scale implementations of our pro-active solution. Another possibility may also include the study of other underlying topologies based on a different neighbor selection process as well as a new threshold selection mechanism based on the number of hops needed to reach the recharge station as opposed to the current distance-based approximation.

References

1. M. I. Afzar, W. Mahmood, and A. H. Akbar. A battery recharge model for wsns using free-space optics (fso). *In Proceedings of the 12th IEEE International Multitopic Conference*, pages 272–277, 2008.
2. M. I. Afzar, W. Mahmood, S. M. Sajid, and S. Seoyong. Optical wireless communication and recharging mechanism of wireless sensor network by using ccrs. *In International Journal of Advance Science and Technology*, 13:59–68, 2009.
3. F. Arwin, K. Samsudin, and A. R. Ramli. Swarm robots long term autonomy using moveable charger. *In Proceedings of the 2009 International Conference on Future Computer and Communication*, pages 127–130, 2009.
4. A. Couture-Beil and R. Vaughan. Adaptive mobile charging stations for multi-robot systems. *In Proceedings of the 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1363–1368, 2009.
5. A. Drenner and N. Papanikolopoulos. Docking station relocation for maximizing longevity of distributed robotic teams. *In Proceedings of the 2006 IEEE International Conference on Robotics and Automation*, pages 2436–2441, 2006.
6. L. Feeney. An energy consumption model for performance analysis of routing protocols for mobile ad hoc networks. *In Mobile Network Applications*, 6:239–249, 2001.
7. L. Feeney and M. Nilsson. Investigating the energy consumption of a wireless network interface in an ad hoc networking environment. *In IEEE Infocom*, 2001.
8. H. Frey, S. Ruhup, and I. Stojmenovic. *Routing in wireless sensor networks*, chapter In Guide to Wireless Sensor Networks. Springer, London, May 2009.
9. W. Heinzelman, A. Chandrakasan, and H. Balakrishnan. An application-specific protocol architecture for wireless microsensor networks. *IEEE Transactions on Wireless Communications*, 1:1536–1276, 2002.
10. E. Kranakis, H. Singh, and J. Urrutia. Compass routing on geometric networks. *In Proceedings of the 11th Canadian Conference on Computational Geometry*, pages 51–54, 1999.
11. X. Li, A. Nayak, and I. Stojmenovic. Exploiting actuator mobility for energy-efficient data collection in delay-tolerant wireless sensor networks. *In Fifth International Conference on Networking and Services ICNS*, pages 216–221, 2009.

12. Y. Litus, R. Vaughan, and P. Zebrowski. The frugal feeding problem: energy-efficient, multi-robot, multi-place rendezvous. *In Proceedings of the 2007 IEEE International Conference on Robotics and Automation*, pages 27–32, 2007.
13. Y. Litus, P. Zebrowski, and R. T. Vaughan. A distributed heuristic for energy-efficient multirobot multiplace rendezvous. *In IEEE Transactions on Robotics*, 25:130–135, 2009.
14. C.-H. Lung, C. Zhou, and Y. Yang. Applying hierarchical agglomerative clustering to wireless sensor networks. *In Proceedings of the International Workshop on Theoretical and Algorithmic Aspects of Sensor and Ad-hoc Networks (WTASA)*, 2007.
15. J. Luo and J. P. Hubaux. Joint mobility and routing for lifetime elongation in wireless sensor networks. *In Proceedings of IEEE INFOCOM*, pages 1735–1746, 2005.
16. Y. Mei, C. Xian, S. Das, Y. Hu, and Y. Lu. Sensor replacement using mobile robots. *Computer Communications*, 30:2615–2626, 2007.
17. F. Michaud and E. Robichaud. Sharing charging stations for long-term activity of autonomous robots. *In Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 3:2746–2751, 2002.
18. M. Perillo, C. Zhao, and W. Heinzelman. On the problem of unbalanced load distribution in wireless sensor networks. *In Global Telecommunications Conference Workshops, 2004. GlobeCom 2004*, pages 74–79, 2004.
19. M. Rahimi, H. Shah, G. Sukhatme, J. Heidemann, and D. Estrin. Studying the feasibility of energy harvesting in a mobile sensor network. *In Proceedings of the IEEE International Conference on Robotics and Automation*, pages 19–24, 2003.
20. S. Roundy, P. Otis, Y. Chee, J. Rabaey, and P. Wright. A 1.9ghz rf transmit beacon using environmentally scavenged energy. *In IEEE International Symposium on Low Power Electricity and Devices*, 2003.
21. M. Sharif, S. Sedighian, and M. Kamali. Recharging sensor nodes using implicit actor coordination in wireless sensor actor networks. *Wireless Sensor Networks*, 2:123–128, 2010.
22. D. Simplot-Ryl, I. Stojmenovic, and J. Wu. Energy efficient backbone construction, broadcasting, and area coverage in sensor networks. *Handbook of Sensor Networks: Algorithms and Architectures*, pages 343–379, 2005.
23. I. Stojmenovic and X. Lin. Power-aware localized routing in wireless networks. *IEEE Transactions on Parallel and Distributed Systems*, 12:1122–1133, 2001.
24. T. Tirta, B. Lau, N. Malhotra, S. Bagchi, L. Z., and Y. Lu. Controlled mobility for efficient data gathering in sensor networks with passively mobile robots. *In IEEE Monograph on Sensor Network Operations.*, 2005.
25. A. Vargas. The omnet++ discrete event simulation system. *In Proceedings of the European Simulation Multi-conference (ESM'2001)*, pages 319–324, 2001.
26. W. Wang, V. Srinivasan, and K. Vikram. Extending the lifetime of wireless sensor networks through mobile relays. *In IEEE/ACM Transactions on Networking*, 16:1108–1120, 2008.
27. J. Warwerla and R. Vaughan. Near-optimal mobile robot recharging with the rate-maximizing forager. *In Proceedings of the 9th European Conference on Artificial Life*, pages 776–785, 2007.
28. O. Younis and S. Fahmy. Heed: A hybrid, energy-efficient, distributed clustering approach for ad hoc sensor networks. *IEEE Transactions on Mobile Computing*, 3:366–379, 2004.