

Neighbour Selection and Sensor Knowledge: Proactive Approach for the Frugal Feeding Problem in Wireless Sensor Networks

Elio Velazquez and Nicola Santoro

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

Abstract. This paper examines new proactive solutions 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 sensors. Our proactive solution reduces the FFP to finding energy-efficient routes in a dynamic Compass Directed Gabriel Graph (CDGG) or Compass Directed Relative Neighbour Graph (CDRNG). The proposed graphs incorporate ideas from forward progress routing and the directionality of compass routing in an energy-aware graph. Navigating the CDGG or CDRNG 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 a shared recharge station. We also examine the impact of critical parameters such as transmission range, number of recharge ports and sensor knowledge for the two proposed graphs.

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, [8, 9] presents 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 how to ensure that the robots reach the rendezvous location without “dying” of energy starvation during the process.

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 sensors are responsible for maintaining the overall health of the network and the service facilities play a passive role. The rendezvous between sensors and facilities should take place at the closest facility’s original position (i.e., static location). 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. Our problem can be seen as the “tanker absorbed” version of the FFP. A similar problem is addressed in [14].

1.1 Related Work

In the FFP, as introduced in [8], specialized robots, also called tankers, have to rendezvous with a team of mobile robots for refueling purposes. The main goal is to minimize the amount of fuel (i.e., energy) required to move the robots and tankers to the rendezvous locations. The problem can have several variants: 1) robot-absorbed case. The rendezvous takes 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. 2) tanker-absorbed case. The rendezvous takes place at the tanker’s location and the robots should move to the tanker’s original location. 3) General case. 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 guarantee that no robots die of energy starvation is an NP-Hard problem (as shown in [8]).

Examples of the robot-absorbed FFP can be found in [1, 2, 10]. In all 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). The simulations results presented in [2] 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 [10] 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 that 50% of the sensors die due to battery depletion.

Reactive vs. proactive strategies for energy restoration in WSN are discussed in [13, 14]. In particular, examples of proactive strategies for the tanker-absorbed FFP can be found in [14] along with the impact of several network parameters such as transmission range, locomotion costs and recharge station role.

1.2 Contributions

This paper emphasizes the use of a proactive approach to solve the Frugal Feeding Problem (FFP) in WSN. 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 Gabriel Graph (CDGG) or Compass Directed Relative Neighbour Graph (CDRNG). We prove that energy-aware mobility strategies built on the CDGG and CDRNG 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 topology, transmission range, number of docking ports and sensor knowledge. This paper also starts a discussion on proactive solutions to the FFP in the presence of obstacles.

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 completely 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 proactive. 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. 4) Our analysis considers the impact of critical network parameters such as neighbour information, transmission range and number of recharge ports.

2 Proactive Alternatives to the Facility-Absorbed FFP

This paper extends some results previously presented in [14]. 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 energy equilibrium where no further losses are reported. 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 other localization methods. 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. The communication environment is contention and error free (i.e., no need to retransmit data) and there is no interference produced by receiving simultaneous radio transmissions (i.e., ideal MAC layer).

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 behaviour 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 higher energy sensors that are closer to the station and eventually making 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) a 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 to the 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.

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 Gabriel Graph (CDGG) or a Compass Directed Relative Neighbour Graph (CDRNG).

Definition 1. Let $G = (V, E)$ be a Unit Disk Graph with vertices V and a set of edges E . A graph $G = (V' \cup F, E')$ with $V' \subseteq V$ and $E' \subseteq E$ is called Compass Directed Gabriel Graph (CDGG) if \forall pair of sensors $s_i, s_j \in V'$ and recharge facility F , the edge $s_i \rightarrow s_j \in E'$ iff the following conditions are satisfied:

1. Unit graph criterion: $d(s_i, s_j) \leq R$ where d denotes the Euclidean distance and R is the transmission range.
2. Proximity criterion: $d(s_j, F) < d(s_i, F)$ and $d(s_i, s_j) < d(s_i, F)$
3. Directionality criterion: $\exists s_{jp}$ such that $\vec{s_j s_{jp}} \cdot \vec{s_i F} = 0$ and $d(s_i, s_{jp}) + d(s_{jp}, F) = d(s_i, F)$
4. Gabriel neighbour criterion: $\nexists s_k \in V'$ such that $d(s_k, \frac{s_i + s_j}{2}) < d(s_i, \frac{s_i + s_j}{2})$

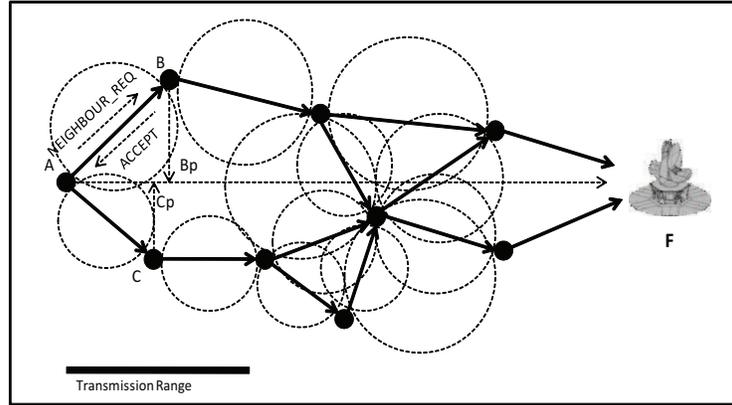


Fig. 1. Compass Directed Gabriel Graph

Definition 2. Let $G = (V, E)$ be a Unit Disk Graph with vertices V and a set of edges E . A graph $G = (V' \cup F, E')$ with $V' \subseteq V$ and $E' \subseteq E$ is called Compass Directed Relative Neighbour Graph (CDRNG) if \forall pair of sensors $s_i, s_j \in V'$ and recharge facility F , the edge $s_i \rightarrow s_j \in E'$ iff the following conditions are satisfied:

1. Unit graph criterion: $d(s_i, s_j) \leq R$ where d denotes the Euclidean distance and R is the transmission range.
2. Proximity criterion: $d(s_j, F) < d(s_i, F)$ and $d(s_i, s_j) < d(s_i, F)$
3. Directionality criterion: $\exists s_{jp}$ such that $\vec{s_j s_{jp}} \cdot \vec{s_i F} = 0$ and $d(s_i, s_{jp}) + d(s_{jp} F) = d(s_i, F)$
4. Relative neighbour criterion: $\nexists s_k \in V'$ such that $d(s_i, s_k) < d(s_i, s_j)$ and $d(s_k, s_j) < d(s_i, s_j)$

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

2.1 Creating the CDGG and CDRNG

Figure 1 shows an example of the proposed CDGG 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 CDGG. The process 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 NEIGHBOUR_REQUEST broadcast message inviting other sensors to participate.
3. Upon receiving a NEIGHBOUR_REQUEST message from sensor A , immediate neighbours verify the neighbouring 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 neighbours 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 NEIGHBOUR_ACCEPT message. Otherwise they send a NEIGHBOUR_DENY message.

Up to this point, the process is the same as the creation of the CDG introduced in [13, 14]. However, to guarantee that only the Gabriel neighbours are selected as graph neighbours, the sensor should implement the following actions:

1. Upon receiving a NEIGHBOUR_ACCEPT message from a potential Gabriel neighbour S' , the receiving sensor S verifies if there is already a graph neighbour in the disc with center $(\frac{S_x+S'_x}{2}, \frac{S_y+S'_y}{2})$ and radius $\frac{d(S,S')}{2}$. If such a neighbour exists, then sensor S sends a NEIGHBOUR_DENY message to S' .
2. If no existing graph neighbour is found in the previous step, this means that sensor S' is in fact a Gabriel neighbour. However, some of the existing graph neighbours could be affected by this newly accepted sensor and they are no longer Gabriel neighbours. If the newly accepted sensor S' falls in the diametric disc between sensor S and one of the existing graph neighbours S_i , the neighbour in question should be excluded by sending it a NEIGHBOUR_DENY message.

The creation of the CDRNG follows the same pattern with only one minor change to verify the relative neighbouring criterion:

1. Upon receiving a NEIGHBOUR_ACCEPT message from a potential relative neighbour S' , the receiving sensor S verifies if there is already a graph neighbour in the Lune created by intercepting the discs with centers in S and S' and radius $d(S, S')$. If such a neighbour exists, then sensor S sends a NEIGHBOUR_DENY message to S' .
2. If no existing graph neighbour is found in the previous step, this means that sensor S' is in fact a relative neighbour. However, some of the existing graph neighbours could be affected by this newly accepted sensor and they are no longer relative neighbours. If the newly accepted sensor S' falls in the Lune between sensor S and one of the existing graph neighbours S_i , the neighbour in question should be excluded by sending it a NEIGHBOUR_DENY message.

The main interactions required for the construction of the CDGG are summarized by Algorithm 1. A detailed description of the CDRNG construction is

omitted for space limitations but it follows the same idea with only minor modifications.

Algorithm 1 GDGG Construction: sensor S and facility F

```

(* In State INIT : *)
send NEIGHBOUR_REQUEST broadcast message
become BATTERY_OK
(* In State BATTERY_OK : *)
if receiving NEIGHBOUR_REQUEST from  $S'$  then
  if  $distance(S, F) < distance(S', F)$  and  $distance(S, S') < distance(S', F)$  and
   $DistancePointToLineIn(S, S', F, distanceToLine)$  then
     $parentList.Add(S')$ 
    send NEIGHBOUR_ACCEPT to  $S'$ 
  end if
end if
if receiving NEIGHBOUR_ACCEPT from  $S'$  then
   $midPoint.X = (S.CoordX + S'.CoordX)/2$ 
   $midPoint.Y = (S.CoordY + S'.CoordY)/2$ 
  while  $i \leq numNeighbours$  do
    if  $S.distance(midPoint) \geq neighbourPositions[i].distance(midPoint)$  then
      send NEIGHBOUR_DENY to  $S'$ 
      become BATTERY_OK
    end if
  end while
  while  $i \leq numNeighbours$  do
     $midPoint.x = (S.CoordX + neighbourPositions[i].CoordX)/2$ 
     $midPoint.y = (S.CoordY + neighbourPositions[i].CoordY)/2$ 
    if  $S.distance(midPoint) \geq S'.distance(midPoint)$  then
      send NEIGHBOUR_DENY to  $neighbour[i]$ 
       $neighbourList.Remove(i)$ 
    end if
  end while
   $rankingPar = d(S, S')$ 
   $neighbourList.Add(S', rankingPar)$ 
   $neighbourList.rank()$ 
end if

```

At the end of this phase each sensor will have two routing tables: one containing its children (i.e., sensors from which NEIGHBOUR_ACCEPT messages were received) with their corresponding ranking and a second table containing its parents (i.e., sensors to which NEIGHBOUR_ACCEPT messages were sent). The routing tables are just partial maps of the network indicating the position of their children and parents. The identity of the sensors in the routing tables is dynamic and will be updated every time a swapping operation occurs. This property, along with a neighboring criteria that incorporates ideas from forward progress and compass routing [11, 7, 6] in an energy-aware unit graph, ensure the following lemma:

Lemma 1. *The proactive solution to the FFP using a CDGG or CDRNG guarantees that all sensors reach the rendezvous location within a finite number of swapping operations.*

Proof. Let $G = (V, E)$ be a CDGG or CDRNG 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

3 Increasing Sensor Knowledge

Another possible enhancement to improve the overall performance of the proactive strategy and help low energy sensors reach the recharge station faster is to add additional information about the energy levels of the 2-hop graph neighbours. Regardless of the topology chosen (i.e., CDG, CDGG, or CDRNG), having the 2-hop neighbouring information combined with the 1-hop greedy strategy should lead to a more energy efficient path selection. To implement this new approach, a series of changes to the existing algorithms is necessary. For example, the neighbouring information stored by each sensor s needs to change to include the tuple $(s_i, E_{S_i}, E_{S_{i_{2hop}}})$ where s_i is the i -th 1-hop neighbour of s . E_{S_i} represents the energy level and $E_{S_{i_{2hop}}}$ represents the average energy levels of the 1-hop graph neighbours of s_i .

The information about existing 1-hop graph neighbours will be appended to the NEIGHBOUR_ACCEPT messages sent during the graph creation phase. When a sensor sends a NEIGHBOUR_ACCEPT message to its parent, the message will now include the average energy level of its existing 1-hop neighbours. This new piece of information will have to be updated once the migration or swapping phase is initiated. Consequently, two swapping sensors will exchange this new piece of information as part of the swapping process. Furthermore, sensors reacting to a SWAP_COMPLETE message will generate a new message NEIGHBOUR_2HOP_UPDATE to inform their parents about the changes of their 2-hop graph neighbours.

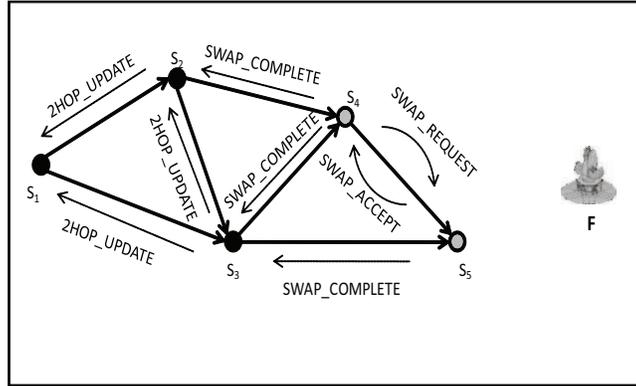


Fig. 2. Sensor swapping with 2-hop neighbours updates.

Let us examine the example shown in Figure 2 to illustrate the new interactions required during a swapping operation. In this example, sensors S_4 and S_5 have agreed to swap positions after the corresponding exchange of SWAP_REQUEST and SWAP_ACCEPT messages. Once the sensors arrive at the location occupied by their swapping partners, both sensors (i.e., S_4 and S_5) will send SWAP_COMPLETE messages to their parents S_2 and S_3 . The SWAP_COMPLETE message received by sensor S_2 contains the tuple $(S_4, E_{S_4}, E_{S_{4_{2hop}}})$. After updating its neighbouring information with the newly received information, S_2 computes the combined energy level of its 1-hop graph neighbours: $E_{S_{2_{2hop}}} = \frac{E_{S_3} + E_{S_4}}{2}$ and sends a new NEIGHBOUR_2HOP_UPDATE $(S_2, E_{S_4}, E_{S_{2_{2hop}}})$ message to its parent S_1 .

It is clear from the previous example that for each successful swapping operation there will be an overhead produced by the new NEIGHBOUR_2HOP_UPDATE messages. The density of the graph, determined by the neighbour selection criteria and the sensor transmission ranges, will have a great impact on how many of these new notification messages are generated. The next section examines the impact of this added knowledge, its relationship with the underlying topology chosen, its potential benefits and possible drawbacks.

4 Simulation 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 [4, 5]. Protocols that rely on periodic probe messages and acknowledgments are considered high cost. It is also noted in the literature that sensors' energy consumption in an idle state can be as large as the energy used when receiving data [5]. On the other hand, the energy used

in transmitting data could be between 30-50% more than the energy needed to receive a packet.

A common consideration for any solution involving mobile entities is how to accurately represent the energy spent when moving from one location to another. Locomotion cost depends on many factors such as the weight of the electronic components, irregularities in the terrain, obstacles, etc. For simplicity, in [8, 14], the weighted Euclidean distance between origin and destination is used as the cost of relocating a robot. In particular, in [14] is observed that the energy required to move their robotic sensors was 54x the energy required to send a packet over the same distance and the energy spent in communications (i.e., send/receive) was 25% more than the battery drain in the idle state.

The simulation scenarios are implemented in Omnet++ [12] along with the mobility framework extension [3]. 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 two 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 and 2) Impact of several variables such as: underlying topology, transmission range, number of recharge sockets/ports and sensor knowledge.

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), it is called a “one-hop run” or “optimal run”. Contrarily, if the final trip is made from any other location, it is called a “panic run” [13]. In all the simulated scenarios, 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 some of the observations found in [14, 4, 5].

4.1 Topology Comparison

This test was designed to determine whether our proactive solution to FFP reaches a state of equilibrium when the new proposed CDGG and CDRNG are used as the underlying topologies for the mobility strategies. The experiment measured the cumulative number of sensor losses until energy equilibrium is reached. Figure 3(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. The sensor transmission range is now 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 $\frac{1}{5}E$ per meter traveled. For all the tests performed on the three different topologies, the mobility strategy selected was the greedy closest-first swapping where a low energy sensor chooses its closest graph neighbour as a swapping partner during its migration towards the recharge station.

As expected, the closest-first swapping strategy on the three topologies chosen (i.e., CDG, CDGG and CDRNG) reached the state of equilibrium. The

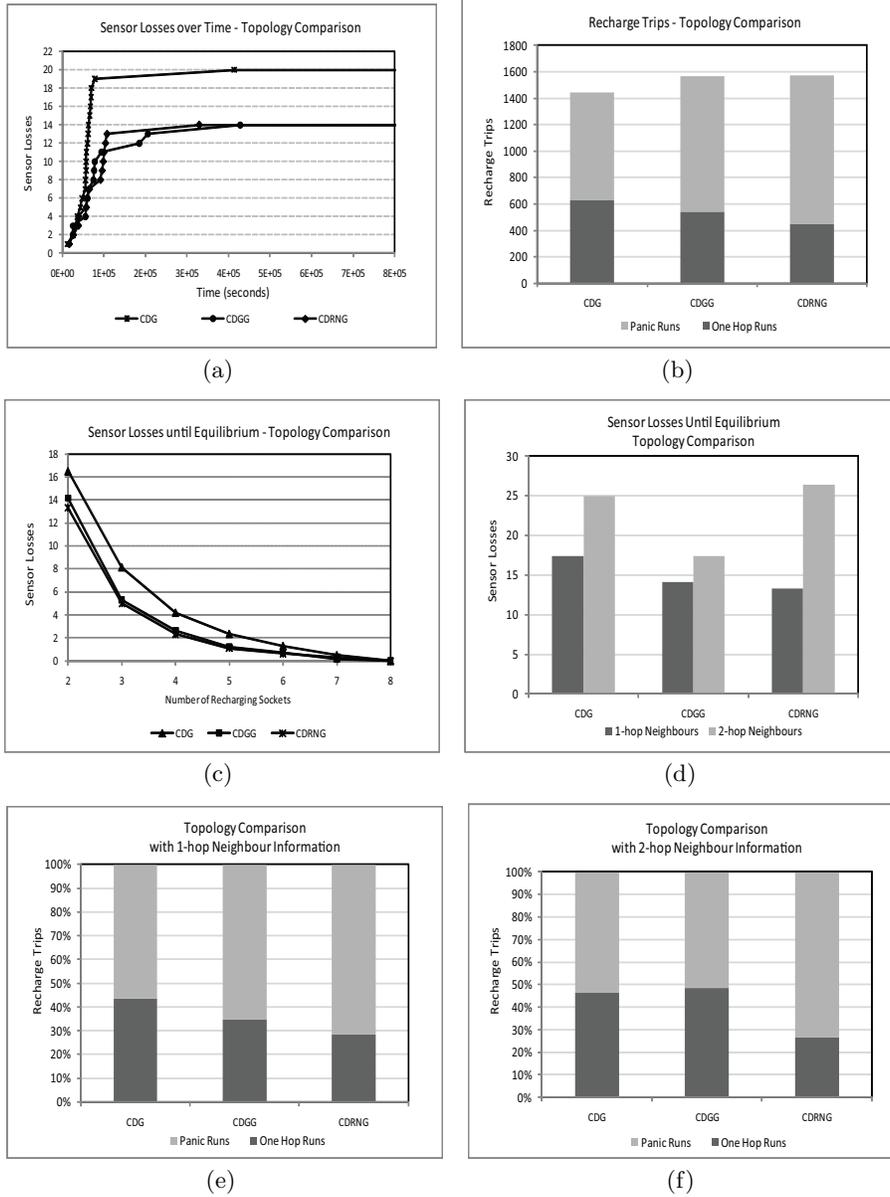


Fig. 3. Simulation Results

CDGG and CDRNG are sub-graphs of the CDG and according to the experimental results presented in [14], even the single path (i.e., single neighbour) approach reached the state of equilibrium. However, the interesting finding is that although the three topologies reached the state of equilibrium at the same

time approximately, the CDGG and CDRNG reported fewer sensor losses due to battery depletion. This is an important observation that implies that fewer but better selected graph neighbours will yield better results if the main goal is to minimize the number of permanent failures due to battery depletion.

Unfortunately, the CDGG and CDRNG did not report any improvements in terms of optimal trips to the recharge station. Figure 3(b) shows the number of recharge trips and breakdown between optimal and panic runs for the three topologies in question. For the CDGG and CDRNG there was a small increase in the number of recharge visits compared to the CDG and a small decrease in the number of optimal runs. This decrease is somehow expected since the number of neighbours for both topologies (i.e., CDGG and CDRNG) is more restrictive than the CDG. Once more, choosing different topologies for the migration strategy exposed a trade-off between permanent coverage holes due to battery depletion and more short-lived temporary holes due to more frequent visits to the facility.

The next part of this test was designed to measure the impact of the recharge sockets on the cumulative number of losses until equilibrium and verify whether the perfect equilibrium can be reached by increasing the number of sockets or docking ports in the recharge station. Figure 3(c) shows the result for this test where the closest-first swapping strategy on the three topologies showed the same progression towards perfect equilibrium. The total number of recharge sockets needed for the perfect equilibrium is the same for the three topologies but the CDGG and CDRNG showed an improvement on the number of sensor losses over the CDG as the number of recharge sockets increased.

4.2 Sensor Knowledge

The goal of this set of tests is to verify the impact of added sensor knowledge, as introduced in 3, and compare it with the 1-hop information greedy strategies on the three topologies (i.e., the CDG proposed in [14] and the CDGG and CDRNG). The network parameters are the same as in the previous tests, with fixed transmission range at 100m. The closest-first swapping strategy is applied on the three topologies (i.e., CDG, CDGG and CDRNG) with information about the energy levels of 1-hop graph neighbours only and 2-hop graph neighbours respectively.

Figure 3(d) shows the number of sensor losses until equilibrium for the three topologies tested with 1-hop neighbour information vs. 2-hop neighbour information. In each case, there was an increase in the number of sensor losses when the migration strategy included the 1-hop neighbour information. When 2-hop information is used, the best performer was the CDGG with losses similar to the 1-hop CDG. This is a rather surprising result, which seems to imply that “knowing more individually” about the network is less useful for the collective effort than “knowing less”. Knowing more in this case has a direct impact on the number of control messages required to maintain the underlying topology in a consistent state. This phenomenon will be more evident as the graph degree increases. The graph maintenance overhead related to keeping 2-hop neighbour

information proved to be crucial to the point that counteracts any possible improvement when compared to keeping 1-hop information only.

The idea of adding extra knowledge to the sensors aimed to improve the path selection strategy and increase the number of optimal runs or 1-hop trips to the recharge station. The simulation results shown in 3(f) confirmed our expectations. Added knowledge had, in fact, a positive impact on the selection of a better energy-efficient migration strategy towards the recharge station. There was some marginal improvement on the number of optimal runs for the CDG and CDRNG with a real improvement for the CDGG. The CDGG proved again to be the best performing topology in terms of cumulative sensor losses until equilibrium and breakdown between panic and optimal runs when using 2-hop neighbour information.

The last test involving the added-knowledge scenario examined the impact of the sensors transmission range on the overall performance. For this test, the closest-first swapping strategy on the CDG with 2-hop neighbour information was implemented on the network of 100:1 sensor/facility ratio with various transmission ranges (e.g., 50m, 100, 200m, 300m).

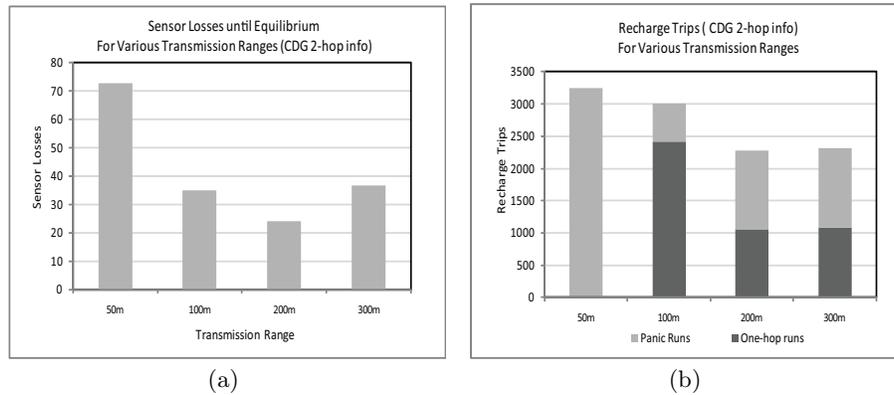


Fig. 4. Simulation Results - Variable Range

Figure 4(a) shows the cumulative number of sensor losses until equilibrium for each range value. The behaviour is very similar to the results for the 1-hop neighbour information scenario presented in [14]. The transmission range of 50m was too restrictive, which means that most of the sensors were isolated and the number of 1-hop and 2-hop neighbours in the CDG was too small to guarantee a gradual approach towards the recharge station. By increasing the transmission range, the number of losses decreased dramatically. However, for the 300m range there was a decline on the overall performance, which is consistent with the 1-hop information scenario.

The number of recharge trips and breakdown between panic and optimal runs is shown in Figure 4(b). Following the same behaviour as in the 1-hop

information scenario, 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 1-hop neighbours. The best breakdown between one-hop and panic runs occurs with the 100m range. However, there are more visits to the recharge location, when compared to the 200m and 300m cases, which reported more balanced results in terms of the number and type of visits to the facility.

5 Conclusions and Future Work

In this work we have enhanced existing proactive strategies to solve the facility-absorbed Frugal Feeding Problem (FFP). Our novel approach proposed the introduction of new underlying topologies with different neighbour selection processes (e.g., Compass Directed Gabriel Graph and Compass Directed Relative Neighbor Graph). The proposed graphs guarantee that sensors will reach the rendezvous location within a finite number of swapping operations with a loop-free migration trajectory. We have also proposed to enhance sensor capabilities and decision making by adding information about energy levels of the 2-hop graph neighbours. All decisions made by the sensors regarding the next swapping operation are based on local knowledge. However, a new look-ahead parameter that includes the combined energy levels of the 2-hop neighbours is taken into account in the selection of the swapping partner.

The simulation results of the modified proactive solution to the FFP show that:

1. For networks of 100:1 sensor/facility ratio, the network survivability rate can be improved by using a CDGG or CDRNG as an underlying topology for the migration strategy.
2. Adding the energy levels of the 2-hop graph neighbours improves a sensor's individual migration strategy towards the facility. There is an increase in the number of optimal trips. However, the number of losses until equilibrium also increases, which results in lower network survivability.
3. If 2-hop neighbour information is available, the proposed CDGG outperforms the other proposed topologies in terms of network survivability at the point of equilibrium and distance traveled to the facility.
4. The transmission range has a positive impact on the network survivability at the point of equilibrium and the number of optimal trips to the facility. However, for higher transmission ranges that result in higher degree graphs, there is a clear negative impact on the key quality indicators (i.e., sensor losses, optimal trips, total number of recharge trips).
5. In general, the simulations exposed several trade-offs between the key variables (i.e., topology, transmission range, locomotion cost, sensor knowledge and station role)

Future enhancements to this work may explore in more detail the proposed proactive strategies in the presence of obstacles and the cost of applying obstacle avoidance strategies. Another possibility may also include the study of other instances of the Frugal Feeding Problem based on the mobility capabilities of

sensors and recharge stations under more realistic MAC layers such as 802.11 CSMA/CA.

References

1. 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.
2. 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.
3. W. Drytkiewicz, S. Sroka, and V. Handziski. A mobility framework for omnet++. *In Proceedings of the 3rd International OMNeT++ Workshop*, 2003.
4. L. Feeney. An energy consumption model for performance analysis of routing protocols for mobile ad hoc networks. *Mobile Network Applications*, 6:239–249, 2001.
5. L. Feeney and M. Nilsson. Investigating the energy consumption of a wireless network interface in an ad hoc networking environment. *In Proceedings of the 20th Annual Joint Conference of the IEEE Computer and Communications Societies (IEEE Infocom 2001)*, 3:1548–1557, 2001.
6. H. Frey, S. Ruhrop, and I. Stojmenovic. Routing in wireless sensor networks. *Chapter 4. Guide to Wireless Sensor Networks*, pages 81–111, 2009.
7. 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.
8. 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.
9. Y. Litus, P. Zebrowski, and R. T. Vaughan. A distributed heuristic for energy-efficient multirobot multiplace rendezvous. *IEEE Transactions on Robotics*, 25:130–135, 2009.
10. M. Sharifi, S. Sedighian, and M. Kamali. Recharging sensor nodes using implicit actor coordination in wireless sensor actor networks. *In Wireless Sensor Network*, 2:123–128, 2010.
11. I. Stojmenovic and X. Lin. Power-aware localized routing in wireless networks. *IEEE Transactions on Parallel and Distributed Systems*, 12:1122–1133, 2001.
12. A. Vargas. The omnet++ discrete event simulation system. *In Proceedings of the European Simulation Multi-conference (ESM'2001)*, pages 319–324, 2001.
13. E. Velazquez and N. Santoro. Mobility-based strategies for energy restoration in wireless sensor networks. *In Proceedings of the 6th International Conference on Mobile Ad-hoc and Sensor Networks (MSN 2010)*, pages 161–168, 2010.
14. E. Velazquez, N. Santoro, and M. Lanthier. Pro-active strategies for the frugal feeding problem in wireless sensor networks. *In Proceedings of the 2nd International ICST Conference on Sensor Systems and Software (S-CUBE 2010)*, 2010.