International Journal of Networking and Computing – www.ijnc.org, ISSN 2185-2847
Volume 10, Number 2, pages 62–83, July 2020

Effective Energy Restoration of Wireless Sensor Networks by a Mobile Robot[1]

Paola Flocchini and Eman Omar
School of Electrical Engineering and Computer Science
University of Ottawa
Ottawa, ON, Canada

and

Nicola Santoro
School of Computer Science
Carleton University,
Ottawa, ON, Canada

Received: February 27, 2020
Accepted: April 1, 2020
Communicated by Akihiro Fujiwara

Abstract

As most existing sensors are powered by batteries, the coverage provided by a sensor network degrades over time and eventually disappears if energy is not restored. A popular approach to energy restoration is to use a *robot* acting as a mobile battery charger/changer. The robot decides where to move next according to a predefined on-line energy restoration strategy. The *effectiveness* of such a strategy depends on the number of nodes it is able to maintain operational at any given time, as well as on for how long a node whose battery is depleted remains non-operational.

The ideal optimal on-line strategy (called `OPTIMAL`) occurs when the robot knows at any time the current status of all sensors, and it computes the best request to satisfy next, based on this information. Although optimal in terms of effectiveness, this centralized strategy would constantly require up-to-date global information; hence its high computational and communication costs make it not feasible.

We consider a drastically different on-line strategy (called `LIC`), which is simple and fully decentralized, uses only local communication, requires no computations, and is highly scalable. In our strategy, the robot visits the sensors in a predefined circular order, moving in a "clockwise" direction and only when aware of a pending request. A sensor whose battery is about to become depleted originates a recharging request and waits for the robot; the request is forwarded according to the circular order in a "counter-clockwise" direction until it reaches either the robot or another sensor waiting for the robot.

We show the perhaps unexpected result that, once the system becomes stable, in most networks the effectiveness of `LIC` is equivalent to that of `OPTIMAL`. In other words, in most cases, in spite of its simplicity and its extremely small (communication and computation) costs, the proposed decentralized strategy is as effective as the optimal centralized one. We augment our theoretical results with experimental analysis, confirming all the analytical results and showing among other things that the system stabilizes very quickly.

[1]A preliminary version of this paper appeared in the proceedings of the Seventh International Symposium on Computer and Networking (CANDAR) 2019.

1 Introduction

1.1 Framework, Problem, and Related Work

Wireless sensor networks are widely employed in a large variety of contexts and applications, mainly to monitor the conditions of the area in which they are deployed. Most existing sensors are powered by batteries whose lifetime is limited; once the battery becomes depleted, the node is no longer sensing and, in absence of redundant coverage, this *sensing hole* creates a *coverage hole* in the monitored area. Indeed, the coverage provided by the network degrades over time and eventually disappears if no action is taken. Extensive research has being carried out on how to address this problem, mainly concentrating on *energy management* strategies, whose goal is to prolong the lifetime of the network and delay the progressive coverage decay by balancing the energy levels among the sensors (e.g., [1, 17]).

A very different line of research has been on *energy restoration*, with the ambitious goal to maintain the network operating perpetually. In this line are proposals to enhance the sensors with (radically different) additional capabilities. For instance, the sensors could be provided with the means to harvest energy from the environment and to convert it to electrical energy, enabling them to recharge their batteries (e.g., [27, 31]). A different direction is to add mobility to the sensors, enabling them to move to recharge facilities deployed throughout the sensing area (e.g., [20, 28]). The drawback of these types of approaches is the increased complexity, and thus cost, of the sensor nodes; this at a time when technology trends are scaling sensors to be smaller and cheaper.

An alternative to adding more complexity to the nodes has been the proposal of using one or more external mobile devices, typically called *robots*, which would go around to restore energy to nodes with (near) depleted energy. The restoration can take place by either *recharging* the depleted batteries or by *replacing* them entirely with fully charged one. Fueled by the recent evolution in wireless power transfer techniques [13], the research on sensor recharging by mobile robots has been quite intensive (e.g., [2,5,16,19,26,34,36]). The alternative of replacement has been considered in the literature (e.g., [21, 32]), albeit with less intensity. For a comparison between these two alternatives see [21].

Notice that the idea of using a mobile robot in sensor networks is not new, as it has been proposed for data gathering and aggregation, for network repairs, as well as for other network maintenance tasks (e.g., [14, 15, 22]). Other examples of these approaches can be found in research papers in robotics (e.g., [3, 6]).

In this paper we are interested in energy restoration by a single mobile robot. Regardless of whether restoration is by recharging or replacement, after servicing a node, the robot must choose where to move next. The algorithm followed by the robots to make this decision, here called *energy restoration strategy*, may prescribe the acquisition of information from the sensors (e.g., energy level, location, etc.) and require possibly complex computations by the robot. Since the decision can be made based solely on current (and older) information, the strategies are necessarily on-line algorithms. In these strategies, a node whose battery is (about to become) depleted issues a request for the robot; the robot moves to service the requests so to optimize some cost parameters, based on the information currently available. Almost all the existing on-demand strategies are however *centralized* (e.g., [9–11, 26]): the information about all the requests is communicated to the robot that then computes where to go next; alternatively, the information is communicated to the base station, which takes the decisions and provides the mobile robot with instructions. In addition to the high communication costs required, the optimization requirements to be met by the decision are typically accompanied by a high computational complexity; these factors imply a difficulty to scale for these strategies. Hence the interest is for *decentralized* strategies.

The existing *decentralized* strategies are [2, 21, 25]; in [21] the concern is to maximize the time until the first interruption of the sensing activity of a single sensor; in [2], which is not on-demand (i.e., the order in which the nodes are recharged is fixed), the total amount of energy that can be put in the system is bounded (i.e. the energy restoration process is limited); in [25] the described technique applies only to linear sensor networks.

1.2 Effectiveness, Costs and Contribution

The fact that some sensors might become inactive is not a problem if an energy restoration strategy is in place that guarantees that every inactive sensor is recharged and becomes operational again. From the point of view of maintaining as much as possible of the network operational at all times, the *effectiveness* of an energy restoration strategy depends on the number of sensors guaranteed to be operational at any given time. We shall call this measure *operational size* or, with an abuse of notation, *coverage*. Note that effectiveness can be equivalently measured by the number of sensing holes in the network at any given time (the smaller the number, the more effective the strategy); interestingly, unless the deployment has $k > 1$ redundant coverage, this measure corresponds to the number of coverage holes at that time.

The other effectiveness measure of interest is the time from the moment a sensor ceases to be operational to the time when the robot arrives to serve it; i.e., for how long a sensing hole lasts. We shall call this measure *disconnection time*.

Associated with each energy restoration strategy are also the computation and communication *costs* incurred when the robot operates in the network according to that strategy.

The effectiveness and the costs of the strategy employed by the robot to service the sensors depends on many factors, a crucial one being the amount of information about the network status available at any given time to the robot.

From the effectiveness point of view, the "ideal" optimal on-line strategy, which we shall call simply `OPTIMAL`, is clearly when the robot knows at any time the current status of all sensors, and it computes on-line the request it must satisfy next so to minimize the number of sensing holes and/or their duration. Even more that any other on-line strategy, in addition to a high computational complexity, `OPTIMAL` would require constantly up-to-date global information; hence the communication costs required to implement it severely limit its feasibility.

In this paper, we propose a drastically different on-line strategy, which we shall call *Local Information and Communication* (`LIC`). In this strategy, the robot visits the sensors in a predefined circular order, moving in a "clockwise" direction when aware of a pending request. A node whose battery is about to become depleted originates a recharging request and waits for the robot; the request is forwarded in a "counter-clockwise" direction until it reaches either the robot or another node waiting for the robot. In other words, each node communicates only locally: with the neighbouring sensors in the circular order (to send or receive a request), or with the robot if currently there (to communicate the presence of a pending request); the robot moves only from one node to the next in the cyclic order, is aware only of whether or not there is a pending request, and has no need of memory or calculation. In contrast to the existing centralized strategies, this simple on-line strategy is fully distributed and *decentralized*, uses only local communication, requires no computations, and it is highly scalable.

1.3 Main Results

We study the effectiveness and costs of `LIC`, both analytically and experimentally in an abstract setting.

Like in every energy restoration strategy, effectiveness depends on two crucial system parameters: the battery life, i.e., the amount of time Δ a fully charged battery lasts under normal operations; and the recharging time, i.e., the amount of time ρ required for the charging/replacement once the robot is at the sensor's site. Let $\mathcal{N}(\Delta, \rho)$ denote the (infinite) class of sensor networks with those specific component characteristics.

We establish several results related to the *stability* of the system under `LIC`, whereas the network is deemed to be in a stable state if the order in which the nodes are charged in a round is the same in every round. In particular, we prove analytically that almost all networks in $\mathcal{N}(\Delta, \rho)$ become stable under `LIC`; when that happens, although the set of operational nodes changes in time, its size remains unchanged; furthermore this value is the same regardless of the initial network size. We also determine the disconnection time for such stable networks, providing a precise characterization of the performance and effectiveness of `LIC`.

We then compare the effectiveness of `LIC` with that of the optimal (but expensive to implement) strategy `OPTIMAL` We show the perhaps unexpected result that, in most networks, the effectiveness of `LIC` is equivalent to that of `OPTIMAL`. In other words, in most cases, in spite of its simplicity and its extremely small (communication and computation) costs, the decentralized strategy is as effective as the ideal optimal one.

We support our theoretical results with experimental analysis, showing that the system stabilizes very quickly and confirming all the theoretical bounds established for coverage size and disconnection time.

2 Model

Let $\mathcal{X} = \{x_0 \ldots, x_{n-1}\}$ be the set of *sensor nodes*, or simply *nodes*, forming the network. Each node has sensory equipment that allows it to monitor its surroundings; it also has provision for wireless communication within a limited range.

Let π be a cyclic order of the nodes; successive nodes in the order (e.g., $x_{\pi(i)}$ and $x_{\pi(i+1)}$, where all operations on the indices are modulo n) are called *neighbours*, and can communicate (possibly by multiple hops).

Normal operations require sensing and occasional communication; both operations consume energy, which is provided by an on-board battery of limited capacity. When the battery is nearly depleted, the node becomes non-operational, thus creating a sensing hole in the network; and the remaining energy is used for local communication necessary for the recharging. In the following, with an abuse of notation, we will say that the battery is depleted when the node becomes non-operational.

Let Δ denote the amount of time it takes for a fully charged battery to become depleted under normal operations. Each node monitors the energy level of its battery and determines whether it is below a fixed threshold. We denote by τ the amount of time, under normal operations, elapsed from the moment the battery falls below the threshold to the time it becomes depleted. In the following, for ease of discussion, we will sometimes refer to Δ as the battery life or capacity, and to τ as the threshold.

A mobile *robot* R is available in the system to recharge/replace the nodes's batteries; once R reaches a node, if the energy level of the node is below the threshold, the robot will restore the energy. We denote by ρ the amount of time it takes for a battery to become replaced/recharged; i.e., if the robot reaches a node x whose battery is below the threshold at time t, x's battery will be fully charged at time $t+\rho$; we assume that x will start being operational as soon as the robot starts restoring its energy. The robot can move from node to node; we denote by $d_{i,j}$ the time it takes the robot to travel from node x_i to x_j. We assume uniform distances among neighbours in the circular order; that is, $d_{\pi(i),\pi(i+1)} = d \geq 1$ $(0 \leq i \leq n-1)$.

We now introduce the measures we use to study the *effectiveness* of energy restoration strategies. Let $\mathcal{S}$ be an energy restoration strategy. The *operational size*, or *coverage*, at time t under $\mathcal{S}$ (denoted by $Coverage(\mathcal{S}, t)$) is the number of operational nodes at that time. Note that the coverage implicitly measures the number $Holes(\mathcal{S}, t) = n - Coverage(\mathcal{S}, t)$ of the sensing holes at time t.

The *disconnection time* for a node x is the amount of time from the moment x becomes inactive to the time when the robot arrives to serve x. Disconnection time is, of course, zero if the node is charged before it becomes inactive. More precisely, the disconnection time for node x at time t under $\mathcal{S}$ (denoted by $Disconnect(\mathcal{S}, t, x)$) is the amount of time x had been inactive when last serviced by the robot before or at time t. This measures indicates for *how long a sensing hole lasts.*

Since the focus of this study is on the effectiveness of the recharging strategies and on their computing and communication costs, we will not address how the robot acquires the means to service the sensor nodes (e.g., by stopping at a recharging station, by extracting it from the environment, etc.) and we assume (as in [30, 36]) that the robot is always capable of doing so.

3 Decentralized Strategy: LIC

We now describe a simple decentralized on-demand strategy, which uses only local information and requires only local communication, hence the name *Local Information and Communication* (LIC).

The proposed strategy assumes the availability of a cyclic ordering π of the nodes (see Figure 1); the construction of such an order is discussed in Section 6. Once each node $x_{\pi(i)}$ has identified its two neighbours $x_{\pi(i-1)}$ and $x_{\pi(i+1)}$, the algorithm can start.

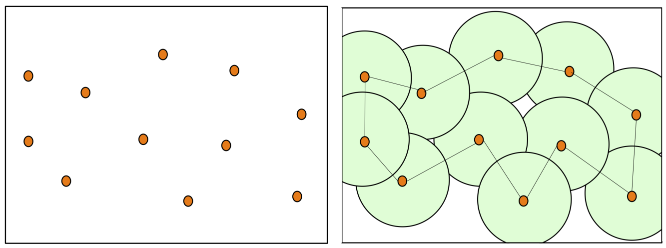

Figure 1: *a*) A deployed sensor network. *b*) A cyclic order π of the nodes.

3.1 Description

Starting from its initial position at an arbitrary node, the robot visits the nodes according to the circular order π, moving in the "counter-clockwise" direction (i.e., from $x_{\pi(i)}$ to $x_{\pi(i+1)}$) when aware of a pending request. A node whose battery is about to become depleted originates a recharging request and waits for the robot; the request is forwarded in the "clockwise" direction (i.e., from $x_{\pi(i)}$ to $x_{\pi(i-1)}$) until it reaches either the robot or another node waiting for the robot, creating in this way a trail to be followed by the robot when it becomes available. Note that a request contains no specific information (e.g., id, location, etc.) about the node issuing or forwarding it. All nodes have the ability to receive and forward a single request even if they are not operational.

Summarizing, (1) each node communicates only locally: with the neighboring nodes in the cyclic order (to send or receive a request), or with the robot if currently there (to communicate the presence of a pending request); (2) the robot moves only from one node to the next in the cyclic order, is aware only of whether or not there is a pending request, and has no need of additional memory or calculation.

Algorithm LIC prescribing the behavior of the nodes and of the robot is shown in Figure 2, using the *State*$\times$*Event*$\rightarrow$*Action* notation. With respect to the protocol, a node x can be in one of two states, Regular or Waiting, and keeps track of whether or not it has received a pending request from its predecessor in the order (Boolean variable $Q(x)$). Initially all nodes are in state Regular, $Q(x) = 0$ for all $x \in \mathcal{X}$, and the robot is at an arbitrary node.

3.2 Properties: Tours and Weakness

A *tour* from node x is defined as the visit of all the nodes by the robot starting from x (and possibly charging it) and ending when arriving again at x. Let $\hat{\Delta} = \Delta - \tau$ denote the amount of time before a fully charged battery falls below the threshold.

Lemma 1. *Let $x_{\pi(i)}$ require recharging both at the beginning and at the end of a tour from it; then also $x_{\pi(i-1)}$ requires recharging when reached by the robot in this tour.*

Proof. Let $x_{\pi(i)}$ be found to be needing recharging at time t_0, fully recharged at time $t_1 = t_0 + \rho$, and found empty again at the end of this tour, at time t_2.

By contradiction, let the previous node $x_{\pi(i-1)}$ be found not needing recharging when visited by the robot in this tour. Let $k \geq 0$ be the number of the nodes (other than $x_{\pi(i)}$) that have been recharged in this tour. In other words, the robot has spent in this tour $k\rho$ time units to charge them

<u>REGULAR</u>
when battery level reaches threshold
  if Robot is not here **then**
    send *request* **to** counter-clockwise neighbour;
  become WAITING;

receiving request
  $Q(x) := 1$;
  send *request* **to** counter-clockwise neighbour;
  become WAITING,

<u>WAITING</u>
receiving request:
  $Q(x) := 1$.

receiving Robot (/*robot arrives here*/)
  if battery level at or below threshold
    `Be Charged`;
  if $Q(x) = 1$ **then**
    $Q(x) := 0$;
    send Robot **to** clockwise neighbour;
  become REGULAR,

Figure 2: Protocol `LIC` executed by node x.

before reaching x_0 at time t_2; that is, $t_2 \geq t_1 + k\rho + nd$. Since $x_{\pi(i)}$ is found needing recharging at time t_2, we must have

$$k\rho + nd \geq \Delta - \tau = \hat{\Delta}.$$

On the other hand, the time elapsed between the time $t' = t_0 - d$ the robot left $x_{\pi(i-1)}$ in the previous tour from $x_{\pi(i)}$ and the time t'' it reached it again in this tour is at least $(k+1)\rho + nd$, since $x_{\pi(i)}$ was recharged in this interval. Since $x_{\pi(i-1)}$ is assumed to be found non needing recharging, we must have

$$(k+1)\rho + nd < \hat{\Delta}$$

a contradiction. □

Analogously,

Lemma 2. *Let $x_{\pi(i)}$ require recharging both at the beginning and at the end of a tour from it; then also $x_{\pi(i+1)}$ will need to be recharged when reached by the robot in the next tour from $x_{\pi(i)}$.*

Proof. Let $x_{\pi(i)}$ be found needing recharging by the robot at time t_0, fully recharged at time $t_1 = t_0 + \rho$, and found needing recharging at the end of the tour, at time t_2. Let $k \geq 0$ be the number of other nodes recharged in this tour. For node $x_{\pi(i)}$ to be needing recharging at time t_2, we must have $k\rho + dn \geq \hat{\Delta}$.

Consider now the state of node $x_{\pi(i+1)}$ when it is reached by the robot once $x_{\pi(i)}$ has been recharged. The time elapsed from the previous visit is greater than or equal to $(k+1)\rho + dn$ (since also $x_{\pi(i)}$ was recharged in this interval). Since $(k+1)\rho + dn = k\rho + \rho + dn > \hat{\Delta}$, it follows that node x_{i+1} needs recharging when it is reached by the robot. □

We say that x is *weak* if there exists a tour from x where x needs recharging both at the beginning and at the end of the tour. Let t_{weak} be the first time when this happens. Lemmas 1 and 2, together, prove the following important property:

Theorem 1. *If there is a weak node then from time t_{weak} every node visited by the robot is found to need recharging. This holds regardless of the threshold τ.*

This, in turns, has important consequences on the size of the coverage of the network:

Theorem 2. *Let $n > m = \lceil \frac{\Delta}{(\rho+d)} \rceil$. If there is a weak node, then there exists a time t, such that, $\forall t' > t, m \leq Coverage(\texttt{LIC}, t') \leq m+1$, and all nodes have the same disconnection time: $\forall x \in \mathcal{X}, Disconnect(\texttt{LIC}, t', x) = (n-1)(\rho+d)+d-\Delta$.*

Proof. Let $x_{\pi(i)}$ be weak; thus, there is a tour where $x_{\pi(i)}$ will be recharged both at the start and the end of that tour; let t be the time when the recharging at the end of that tour will be completed. By Theorem 1, from this time on the robot finds only nodes needing recharging; since it takes ρ time units for the robot to recharge a node and d time units to move to the next node, by time $t' = t + \Delta$ the robot has recharged the consecutive nodes $x_{\pi(i+1)}, x_{\pi(i+2)}, ..., x_{\pi(i+m-1)}$ where $m = \lceil \frac{\Delta}{(\rho+d)} \rceil$ and all operations on the indices are modulo n; if Δ is not a multiple of $\rho + d$, then it is currently recharging $x_{\pi(i+m)}$, otherwise also $x_{\pi(i+m)}$ has been fully recharged. But, at this time $t' = t + \Delta$, $x_{\pi(i)}$'s battery is totally depleted, and so obviously is the battery of all the nodes after $x_{\pi(i+m)}$ up to and including $x_{\pi(i)}$. This means that, at time $t_0 = t + m(\rho + d)$, exactly $n - m$ nodes have their battery completely empty; hence $Coverage(\texttt{LIC}, t_0) = m$.

Observe that, when $x_{\pi(i+m+1)}$ is reached and fully recharged, at time $t_1 = t_0 + (\rho + d)$, $x_{\pi(i+1)}$'s battery is depleted; that is $Coverage(\texttt{LIC}, t_1) = m$. More generally, when $x_{\pi(i+m+j)}$ is reached and fully recharged, at time $t_j = t_0 + j(\rho+d)$, $x_{\pi(i+j)}$'s battery is depleted; that is, $Coverage(\texttt{LIC}, t_j) = m$. On the other hand, at any time $t_j < t' < t_{j+1}$ $x_{\pi(i+j)}$'s battery might not yet be depleted; that is, $Coverage(\texttt{LIC}, t') \leq m + 1$.

Since after time t the robot keeps charging every node it encounters, it will spend $n(\rho + d)$ time units to complete a tour. During that time, each node is not disconnected for $\rho + \Delta$ time units. Therefore, every node will have disconnection time $(n-1)(\rho + d) + d - \Delta$. □

3.3 Properties: Rounds and Stability

Let us call *round* from node x any sequence of consecutive tours from x where at the beginning of the first tour and at the end of the last x's battery needs to be recharged, and in all others it does not. Clearly a round from x might include several tours from x.

We will denote by $r_j(x)$ the j-th round from x, $j \geq 1$; when no ambiguity arises, we will indicate a round from x simply by $r(x)$.

Let $\sigma_j(x)$ denote the ordered sequence of the nodes charged during $r_j(x)$. We say that the network is *stable* if $\exists j \geq 1$ such that $\forall x \in \mathcal{X}, \forall j' > j$, $\sigma_j(x) = \sigma_{j'}(x)$. That is, in a stable network, the order in which the nodes are *charged* is the same in every round; hence, in a stable network, we can omit the indication of the round and denote $\sigma_j(x)$ simply as $\sigma(x)$.

When a network is (or has become) stable, it enjoys some obvious properties. Among them:

Lemma 3. *Let the network be stable. Then*
(i) $\forall x \in \mathcal{X}, \sigma(x)$ is a permutation of the elements of $\mathcal{X}$;
(ii) $\forall x, y \in \mathcal{X}$, $\sigma(x)$ is a cyclic shift of $\sigma(y)$;
(iii) every round from any node is composed of the same number of tours.

An important property of a network once it becomes stable under `LIC` is the following.

Theorem 3. *If the network is stable under* `LIC` *and $n > 2\frac{\Delta}{\rho} + 1$, each round is composed by a single tour.*

Proof. Let the network be stable; then, by Lemma 3 (iii), every round from any node x is composed of the same number s of tours. We want to show that $s = 1$ if $n > 2\frac{\Delta}{\rho} + 1$.

By contradiction, let $s > 1$. Consider a round from node x and let $f(x, s)$ be the number of nodes charged in the last tour of this round. We will consider three cases depending on the value of $f(x, s)$.

Case 1 : $f(x,s) < \frac{n}{2}$. In this case, the number of nodes charged in the first $s-1$ tours is $k = n - f(x,s) \geq \lfloor \frac{n}{2} \rfloor + 1$.

Consider the amount of time T elapsed from the moment x has been charged at the beginning of the round, to the beginning of the last tour of this round.

By definition, and since $k \geq \lfloor \frac{n}{2} \rfloor + 1$ and $s > 1$, we have

$$T = (k-1)\rho + (s-1)d\ n \geq \lfloor \frac{n}{2} \rfloor \rho + (s-1)d\ n \geq \lfloor \frac{n}{2} \rfloor \rho + d\ n.$$

Since, by definition of round, x is found by the robot not to need recharging at this time, we have $\Delta - \tau = \hat{\Delta} \geq T$, that is:

$$\Delta \geq \hat{\Delta} \geq T \geq \lfloor \frac{n}{2} \rfloor\ \rho + d\ n > \lfloor \frac{n}{2} \rfloor\ (\rho + 1)\ .$$

But $n > 2\frac{\Delta}{\rho} + 1$ by hypothesis; that is, $\frac{(n-1)}{2}\rho > \Delta$; a contradiction.

Case 2 : $f(x,s) > \frac{n}{2}$. In this case, there must exist a node y such that $f(y,s) < \frac{n}{2}$. By considering the round from y (instead than from x), by Case 1 the contradiction occurs.

Case 3 : $f(x,s) = \frac{n}{2}$. If $s > 2$ then there must exist a node y such that $f(y,s) < \frac{n}{2}$; by considering the round from y (instead than from x), by Case 1 the contradiction occurs.

Finally, let $s = 2$. In this case, n is even, the round $r(x)$ is composed of two tours, and $f(x,1) = f(x,2) = \frac{n}{2}$. Note that the nodes charged in the second tour of $r(x)$ are the complement of the ones charged in the first tour.

Consider now the node y, next in the cycle, visited by the charger right after x. We have two cases (see Figure 3) depending on whether or not y needs to be recharged at this time. Case $(3a)$: y needs to be charged. In this case, consider the round $r(y)$ starting after charging y (see Figure 3 (3a), bottom). Round $r(y)$ must also be composed of two tours each containing exactly $\frac{n}{2}$ nodes in need of charge (i.e., we must have $f(y,1) = f(y,2) = \frac{n}{2}$), otherwise a contradiction would arise because of the previous reasoning applied to y. However, the nodes in need of charge in the first tour of $r(y)$ are the same as the ones in the first tour of $r(x)$ except for y itself; thus $f(y,1) = \frac{n}{2} - 1$, a contradiction. Case $(3b)$: y does not need to be charged. In this case, consider the round $r(y)$ starting from y the last time it was charged before time $t_s(x)$. By definition, the nodes in need of charge in the first tour of $r(y)$ are the complement of the ones charged in the first tour of $r(x,1)$ excluding y, thus $f(y,1) = \frac{n}{2} - 1$, a contradiction. □

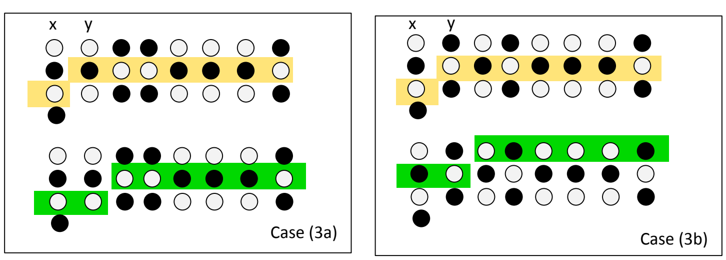

Figure 3: Cases (3a) and (3b) of Theorem 3 in a stable network with $n = 8$. Each row corresponds to a tour starting from x; black (white) circles represent nodes charged (not charged) in that tour. A round $r(x)$ consists of two consecutive tours starting after a black x. Highlighted in the top (resp. bottom) is the first tour of $r(x)$ (resp. $r(y)$) in the two cases.

4 Stability, Effectiveness, and Optimality

4.1 Stability and Effectiveness

All the analytical properties we have established on `LIC` hold if and when the network becomes stable. In the following, we prove that stability under `LIC` is inevitably achieved in almost all networks, establish an upper bound on the number of rounds before that happens, as well as tight bounds on the effectiveness of `LIC`.

4.1.1 Stability of `LIC`

In the following, we prove that, in all networks with $n > \frac{2\Delta}{\rho} + 1$, stability under `LIC` is always achieved. To do so, we need some additional terminology and lemmas.

Given a round $r(x)$ from x, let $tour_1(r(x)), tour_2(r(x)), \ldots, tour_s(r(x))$ denote the tours it is composed of, and let $k(x,i)$ be the number of nodes charged in $tour_i(r(x))$. In particular, let $k(x) = \sum_{i=1}^{s} k(x,i)$ and $\overline{k}(x) = k(x) - k(x,s)$. In the following, let $n > 2\frac{\Delta}{\rho} + 1$.

Lemma 4. *If there exists a node x starting a round $r(x)$ at time t with $s = 1$, then the network is stable from time t.*

Proof. A node x starting a round $r(x)$ at time t with $s = 1$ is, by definition, a weak node. From Theorem 1, we know that from time t every node visited by the robot is found to need recharging, that is, after this time, the order in which the nodes are charged is the same in every round and thus the network is stable by definition.

□

We next show that a round from any node which is composed by more than one tour (i.e. with $s > 1$) have the property that, during the first $s - 1$ tours of the round, less than half of the nodes are charged.

Lemma 5. *There exist no node x completing a $r(x)$ with $s > 1$ tours with $\overline{k}(x) \geq \frac{n}{2}$.*

Proof. By contradiction. Consider a node x starting a round composed by $s > 1$ tours, such that its first $s-1$ tours contain at least $\frac{n}{2}$ recharges; i.e., $\overline{k}(x) \geq \frac{n}{2}$. Since x is found in no need of charge after $s - 1$ tours, we have that $\rho\overline{k}(x) + (s-1)nd < \hat{\Delta}$. However, by hypothesis, $\rho\overline{k}(x) + (s-1)nd > \rho\frac{n}{2}$, which implies $\rho\frac{n}{2} < \hat{\Delta}$, and thus

$$n < 2\frac{\hat{\Delta}}{\rho} < 2\frac{\Delta}{\rho} + 1,$$

a contradiction. □

In other words, a necessary condition for having only rounds composed by more than one tour is that $\overline{k}(x) < \frac{n}{2}$ for all nodes and all rounds.

Lemma 6. *If there exists a node x starting round $r(x)$ at time t with $k(x,s) \geq \frac{n}{2}$, then there exists a node x' starting $round(x')$ at some time $t' > t$ such that $round(x')$ is composed by a single tour (i.e., x' is weak).*

Proof. Consider the last node x' charged in $tour_{s-1}(x)$ and $round(x')$ starting after the charge of x'. Assume by contradiction that $round(x')$ is composed by more than a single tour. Since $round(x')$ fully includes the $k(x,s)$ nodes charged in $tour_s(x)$ and since, by hypothesis, $k(x,s) \geq \frac{n}{2}$ we have that $\overline{k(x')} > \frac{n}{2}$, which is impossible by Lemma 5. □

Lemma 7. *If a node is not charged in an entire round it will be charged in the first tour of the next round.*

Proof. Let y not be charged in round $r(x)$ from x, started at time t_1 and ended at time t_2. Since x needed to be recharged at time t_2, then $t_2 - (t_1 + \rho) \geq \hat{\Delta}$. But since y was charged before time t_1, this means that at time t_2 also its battery life is below the threshold; hence it will be charged as soon as it is reached by the robot in the first tour. □

We can now prove the main stability result.

Theorem 4. *Let $n > \frac{2\Delta}{\rho} + 1$. Then, under the `LIC` strategy, the network becomes stable within finite time.*

Proof. Consider the first node x that is recharged twice by the robot since the beginning. In other words, the first round $r_1(x)$ from x is the first to be completed among all the first rounds from all nodes. Let t_1 and t_2 be the time when $r_1(x)$ starts and ends, respectively. Let $r_1(x)$ be composed of s tours.

If $s = 1$, then the network is stable from time t_1 by Lemma 4, proving the Theorem.

Let then $s > 1$. In this case, by Lemma 5, the number of nodes, other than x, recharged by the robot in the first $s-1$ tours of $r_1(x)$ is $\overline{k_1}(x) < \frac{n}{2}$. Consider the total number $k_1(x)$ of nodes, other than x, recharged by the robot in $r_1(x)$.

If $k_1(x) > \frac{n}{2}$, let x' be the last node charged in the $(s-1)$-th tour of $r_1(x)$; by Lemma 6, its first round $r_1(x')$, started at time $t' < t_2$, is composed by a single tour; hence, by Lemma 4, the network is stable from time t', proving the Theorem.

Finally, consider the case $k_1(x) \leq \frac{n}{2}$. Clearly, in this case, the number $\bar{k} = n - k_1(x)$ of nodes that are not charged at all during $r_1(x)$ is $\bar{k} \geq \frac{n}{2}$. By Lemma 7, all of them will be charged in the first tour of the next round $r_2(x)$ from x. This implies that, if $r_2(x)$ had more than one tour, then $\overline{k}_2(x) > \frac{n}{2}$ contradicting Lemma 5; hence the network is stable from time t_2, proving the Theorem. □

4.1.2 Effectiveness of `LIC`

By bringing together the result of Theorem 4 with the observations on weakness and stability made in Sections 3.2 and 3.3, we can immediately establish tight bounds on the effectiveness of `LIC`.

Theorem 5. *Let $n > \frac{2\Delta}{\rho} + 1$. Then, under the* `LIC` *strategy, there is a time t such that, for all $t' \geq t$, we have*

$$m \leq Coverage(\texttt{LIC}, t') \leq m + 1,$$

where $m = \lceil \frac{\Delta}{(\rho+d)} \rceil$; moreover, for all $x \in \mathcal{X}$

$$Disconnect(\texttt{LIC}, t', x) = (n-1)(\rho + d) + d - \Delta.$$

Proof. By Theorem 4, the network will become stable within the first two rounds. By Theorem 3, it follows that, from time t_{stable}, each round is composed by a single tour; hence, by definition, every node is weak. The bounds on coverage and disconnection time then follow from Theorem 2. □

4.1.3 Round Bounds for Stability

We have shown that all networks with $n > \frac{2\Delta}{\rho} + 1$ become stable under `LIC`. In this subsection we investigate how long it takes for stability to be reached, and establish an upper bound on the number S of tours within which stability will always occur.

Consider the first node x that is recharged twice by the robot since the beginning. In other words, the first round $r_1(x)$ from x is the first to be completed among all the first rounds from all nodes. Let $r_1(x)$ be composed of s tours; let us see how large can s be. From the last case considered in the proof of Theorem 4, it follows that the network becomes stable after at most one more round; that is $S \leq s+1$. Consider the number $\overline{k}_1(x)$ of nodes, other than x, recharged by the robot in the first $s-1$ tours of $r_1(x)$; then obviously, $s \leq \overline{k}_1(x) + 1$. Since by Lemma 5, $\overline{k}_1(x) < \frac{n}{2}$ when $s > 1$, we have $s \leq \frac{n}{2}$. Hence, we have the first upper bound

$$S \leq \frac{n}{2} + 1$$

This bound can be refined. In addition to x, at least one node is recharged in each tour of $r_1(x)$. However, since by Lemma 5 $\overline{k}_1(x) < \frac{n}{2}$, $k_1(x) - \overline{k}_1(x) > \frac{n}{2}$ nodes will not be recharged in the first $s-1$ tours. Among them, let y be the one closest to x in the cyclic order. Since y has not been recharged in the first $s-1$ tours, it means that its energy level when visited by the robot in the $(s-1)$-th tour was high enough. This in spite of the fact that, from time t_1 when this round started, after x is fully recharged at time $t_1 + \rho$, at least one node is recharged in each tour of $r_1(x)$. In other words, $\Delta - \rho > (s-1)\rho + (s-1)nd$; that is, $\Delta > s(\rho + nd) - nd$; hence $\frac{\Delta}{(\rho+nd)} + 1 > s$ which gives us a better upper bound

$$S < \frac{\Delta}{(\rho + nd)} + 2$$

The experimental evidence, to be shown in Section 5, indicates that this bound is quite far from the actual value of S.

4.2 `LIC` versus `OPTIMAL`

We are going to compare the effectiveness of `LIC` with that of the optimal on-line strategy `OPTIMAL`.

In `OPTIMAL`, each request message is sent by the node to the charger; the robot processes all the current request messages, and it computes which request to satisfy next so to minimize the number of sensing holes and their duration. We are actually going to consider an *ideal* cost settings for `OPTIMAL`: every request from every node reaches the robot directly, regardless of its current location; the robot can reach any node from any node in the same amount d of time, regardless of its distance; and the robot's processing time is negligible regardless of the complexity of the computation.

Notice that the behaviour of the robot under `OPTIMAL` in this setting is easy to describe: the robot just processes and services the request messages in the order they arrive; if two or more requests arrive at the same time, ties are broken by Ids.

The effectiveness of `OPTIMAL` is simple to derive for most networks:

Theorem 6. *If $n > (\Delta + \rho)/(\rho + d)$ then, under the* `OPTIMAL` *strategy, there exists a time t such that, for all $t' > t$*

$$\lceil \tfrac{\Delta}{(\rho+d)} \rceil \leq Coverage(\texttt{OPTIMAL}, t') \leq \lceil \tfrac{\Delta}{(\rho+d)} \rceil + 1;$$

moreover, for all $x \in \mathcal{X}$

$$Disconnect(\texttt{OPTIMAL}, t', x) = (n-1)(\rho + d) + d - \Delta.$$

Proof. Let $x_0, x_1, \ldots x_{n-1}$ be the nodes ordered by their initial battery level, where nodes with the same level are ordered by their Id. Clearly this is the order in which the requests are considered by the robot. Let $t_{i,1}$ be the time when x_i is fully charged for the first time; thus, $t_{i,1} + \rho + d \leq t_{i+1,1}$. This implies that the time $t_{i,j}$ when the node x_i will been fully charged for the j-th time is $t_{i,j} + \rho + d \leq t_{i+1,j}$. In other words, the nodes will be always serviced according to the initial order.

After fully charging node x_i for the j-th time at time $t_{i,j}$, the robot will have fully charged all other nodes after at least $(\rho + d)(n-1)$ time units. By hypothesis, $n \geq \frac{\Delta+\rho+d}{\rho+d}$; thus, $(\rho + d)(n-1) \geq \hat{\Delta}$. In this case, by time $t_{i,j} + (\rho + d)(n-1)$, the $(j+1)$-th request by x_i is already in the queue and thus the robot goes to service x_i next. In other words, after a full round, the robot continuously charges the nodes in the same order without ever stopping.

The coverage size in this case is easy to compute. Let X be the total number of nodes that are alive at any given time. This means that $\rho(X-1) + (X-1)d < \Delta$, but $\rho X + Xd \geq \Delta$ and thus $\frac{\Delta}{(\rho+d)} < X < \frac{\Delta+\rho+d}{(\rho+d)}$.

Since after time $t_{0,2}$ the robot keeps charging every node it encounters, it will spend $(n-1)(\rho + d) + d$ time units to complete a round; hence, every node will have disconnection time $(n-1)(\rho + d) + d - \Delta$. $\square$

This theorem has an immediate very strong consequence for the effectiveness of `LIC`.

Theorem 7. *If $n > \frac{2\Delta}{\rho} + 1$ then there exists a time t such that for all $t' > t$*

$$Coverage(\texttt{LIC}, t') = Coverage(\texttt{OPTIMAL}, t');$$

moreover, for all $x \in \mathcal{X}$

$$Disconnect(\texttt{LIC}, t', x) = Disconnect(\texttt{OPTIMAL}, t', x)$$

Proof. Since $\frac{2\Delta}{\rho}+1 > \frac{\Delta+\rho}{\rho+d}$, the claim follows directly from Theorems 5 and 6. □

In other words, for all networks with $n > \frac{2\Delta}{\rho}+1$, the recharging strategy `LIC`, with its low communication and computations costs, performs as well as the optimal strategy.

5 Experimental Analysis

The results of the previous Section describe the behaviour of the network once it stabilizes in time. We run extensive simulation to determine the stability of the network under `LIC` and to observe its behaviour in terms of Coverage Size and Disconnection Time (already studied theoretically).

5.1 Simulation Environment

The experiments were implemented in the simulator discrete-event MAS toolkit MASON [18].

The variable parameters involved in the experiments are: the number of nodes n, the battery life Δ, the charging time ρ, and the travel time d from a node to the next.

The following table shows the values considered for each parameter, where the temporal values are all in the same scale:

Parameters	*Values*
Number of nodes n	100, 200, 300, 400, 500
Battery Lifetime Δ	2000, 3000, 4000
Threshold τ	30% of Δ
Charging Time ρ	1, 10, 20, 30, 40, 50
Travel Time d	1, 5, 10, 20, 30

Each node x has initially an amount of energy level chosen uniformly at random in the range [τ, Δ]; the charger's initial placement is at a node chosen uniformly at random.

For each combination of the values of the parameters, we have executed 20 executions and computed the average coverage size and disconnection time. To detect stability we have also maintained the information about the charging order of the nodes under `LIC`.

5.2 Stability

Starting with arbitrary initial charge levels, the experimental results show that, for all values of the parameters, the charging order becomes periodic and the network stable; moreover stabilization occurs within two rounds. In other words, *under* `LIC` *the network becomes stable within two rounds.* See Figure 4 where stability is shown for some choices of the parameters; the results for all the other parameters' combinations are consistent with these. Note that, for smaller capacities, as well as for larger networks, stabilization occurs even sooner, within one round.

Once the network stabilizes, the theoretical bounds on coverage size and disconnection time established analytically (Theorem 5) hold; indeed, the simulation results of the next section confirm all these bounds.

5.3 Disconnection Time

We now study the impact of varying the various network parameters on Disconnection Time. The simulations show expected results without revealing any surprises. Generally, Disconnection Time increases with the increase of the number of nodes, of charging time and of travel time between successive nodes, while it decreases with the increase of capacity.

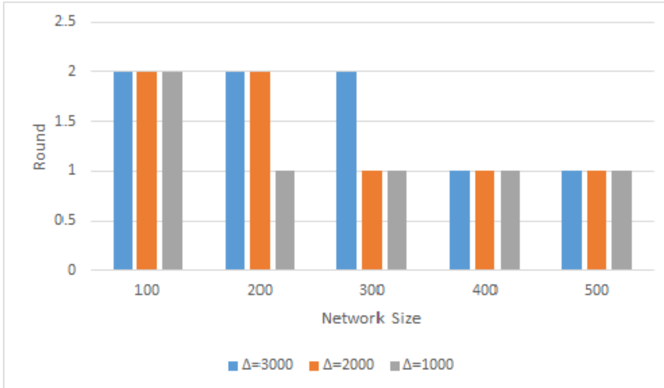

Figure 4: stabilization n=300, d=1

5.3.1 Effects of Charging Time ρ.

In order to determine the impact of the charging time ρ on the effectiveness of the proposed algorithm, we start by analyzing the effects of various charging time on disconnection time for different network sizes using $n = 100, 200, 300, 400, 500$ nodes. In these scenarios, we fix the node capacity $\Delta = 2000$ and we consider unitary travel time between neighboring nodes (the case of $d \neq 1$ will be treated separately). The results obtained varying Δ are consistent with the ones displayed here and are not shown. In all cases the results are as one would have expected.

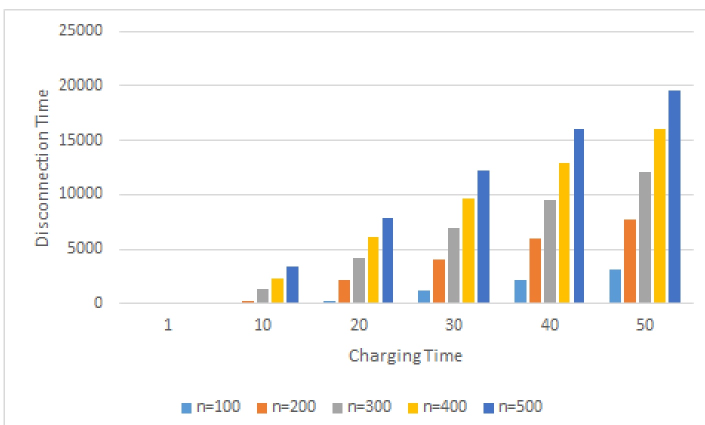

Figure 5: Effect of Charging Time ρ on Disconnection Time; Δ=2000, d=1

The graphs in Figure 5 shows the trend in the disconnection time for the considered network sizes. In general, we can see a similar pattern in the graph: for a given number of nodes, the disconnection time increases as the charging time increases because the waiting time for the robot increases.

5.3.2 Effects of Node Capacity Δ.

We now analyze the effect of node capacity Δ. In the graphs below we show the impact of Δ by fixing n and varying ρ, as well as by fixing ρ and varying n. Also here, the travel time between nodes is considered unitary (the case of $d \neq 1$ will be treated separately). Figure 6 shows the impact on disconnection Time (a), (b): when $n = 300$ nodes while ρ varies. The results reveal no surprises: for a given network size and charging time, the disconnection time decreases as the node capacity increases.

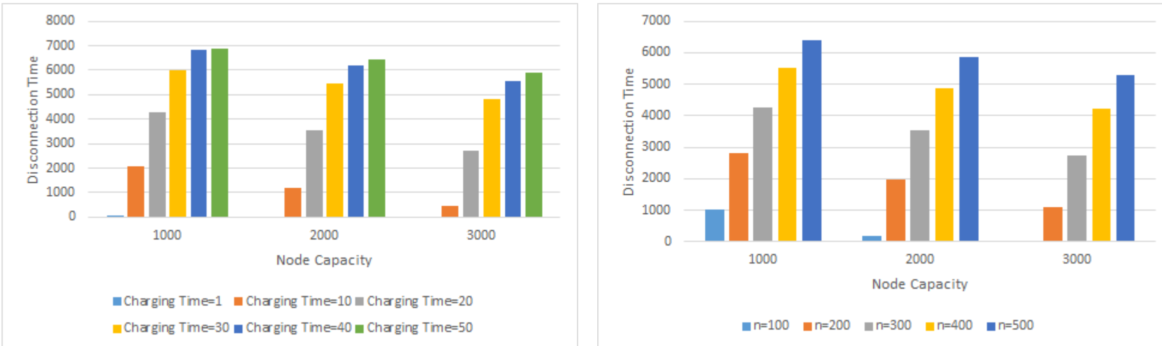

(a) Effect of Node Capacity on Disconnection Time. n= 300

(b) Effect of Node Capacity on Disconnection Time.ρ=20

Figure 6: Effect of Node Capacity Δ on Disconnection Time

5.3.3 Effects of Network Size n.

We now observe the effect of network size on disconnection time, fixing the node capacity $\Delta = 2000$ and travel time $d = 1$, and varying the value of the charging time ρ. As Figure 7 shows, disconnection

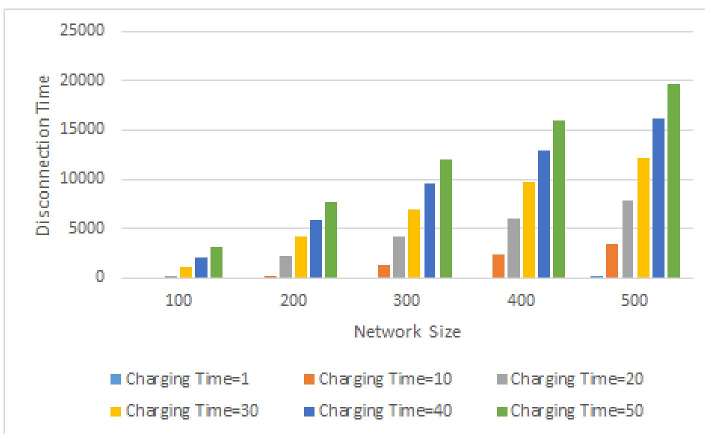

Figure 7: Effect of Network Size. Δ=2000, d=1

time increases, as expected, while the network size increases, and they do so also when varying ρ.

5.3.4 Effects of Travel Time.

We now study the impact of the time it takes the robot to go from a node to the next in the cyclic order on disconnection time. In the experiments, we set node capacity $\Delta = 2000$ unit of energy, and we vary the travel time between nodes to be uniform, but not necessarily unitary: we consider $d = 1, 5, 10$ units. We consider various network sizes $n = 100, 200, 300, 400, 500$, and various charging time $\rho = 1, 10, 20, 30, 40, 50$. Figures 8 depicts the effect of the travel time on the disconnection time. We conclude that as the travel time between nodes increases, the disconnection time increases too. In general, most nodes get charged before their battery becomes empty, which means a lower disconnecting time. As the number of inactive nodes increases due to the increase in travel time, the disconnecting time increases.

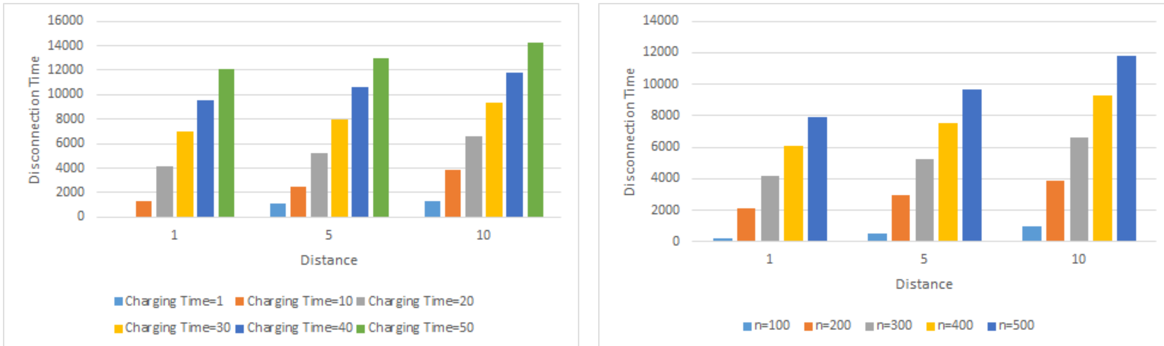

(a) Effect of Travel time on Disconnection Time.n=300 (b) Effect of Travel time on Disconnection Time.ρ=20

Figure 8: Effect of Travel time on Disconnection Time. Δ=2000

5.4 Coverage Size

While the impact on disconnection time is as expected, the observation of coverage size has revealed much more interesting behaviors.

5.4.1 Effects of Charging Time ρ.

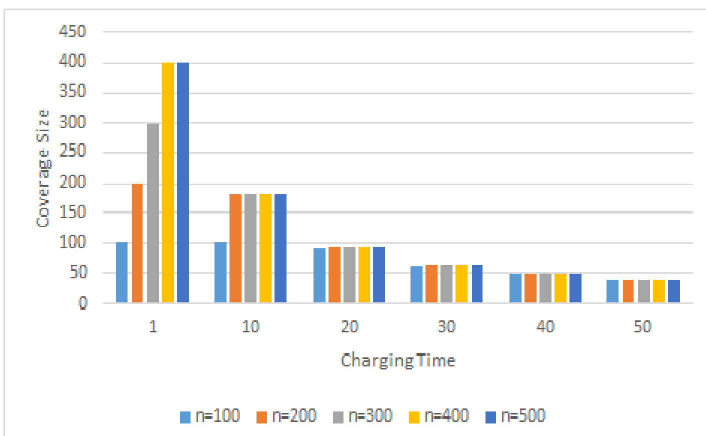

Figure 9: Effect of Charging Time on Coverage Size. Δ=2000, d=1

From Figure 9 we observe that when $\rho = 1$ the increase in size corresponds to an increase of coverage size up to $n = 400$, at which point the coverage size seems to stabilize to a constant value. The same behavior can be observed for larger ρ, where however the stabilization occurs earlier; in fact, for $\rho \geq 10$ we observe it when $n = 200$. This behavior is quite interesting and will be discussed further later. From the graph, we can also see that coverage size decreases varying ρ, when observing a fixed value of n, which is to be expected.

5.4.2 Effects of Node Capacity Δ

Figure 10 (a) shows that, as one would expect, when node capacity increases, also the coverage size increases (when n is fixed). On the other hand, fixing a given capacity, the increase in the network size correspond to a decrease of coverage size. Interestingly, from Figure 10 (b) we observe again that the coverage size does not change when fixing a given node capacity, even when n changes. In other words, the node capacity is independent of the network size and confirms the theoretical derived in Theorem 5, even for n smaller that $\frac{2\Delta}{\rho} + 1$.

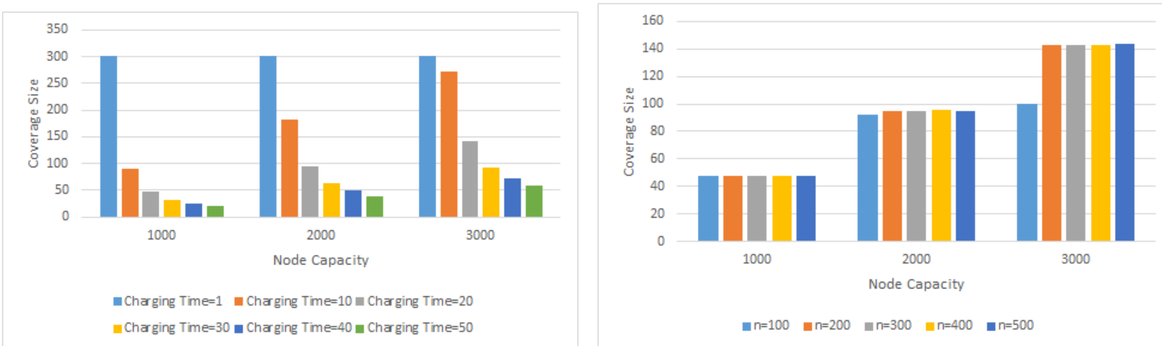

(a) Effect of Node Capacity on Coverage Size. $n = 300$ (b) Effect of Node Capacity on Coverage Size. $\rho = 20$

Figure 10: Effect of Node Capacity on Coverage Size

5.4.3 Effects of Network Size n.

The observation just made can be clearly seen also in Figure 11, where we note that the coverage size increases with the network size but eventually stabilizes. The stabilization occurs earlier for smaller values of charging time ρ still confirming the theoretical bounds established in the previous Section.

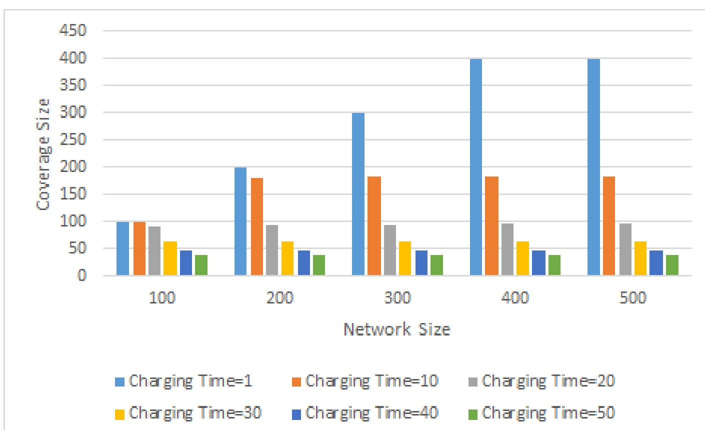

Figure 11: Effect of Network Size on Coverage Size. Δ=2000, d=1

5.4.4 Effects of Travel Time.

When the travel time between two nodes is not unitary, but still constant, we make observations that are consistent with what already seen (see Figure 12), which makes us conclude that, as long as the travel time between two nodes is the same, the actual value does not influence the general behavior of the strategy.

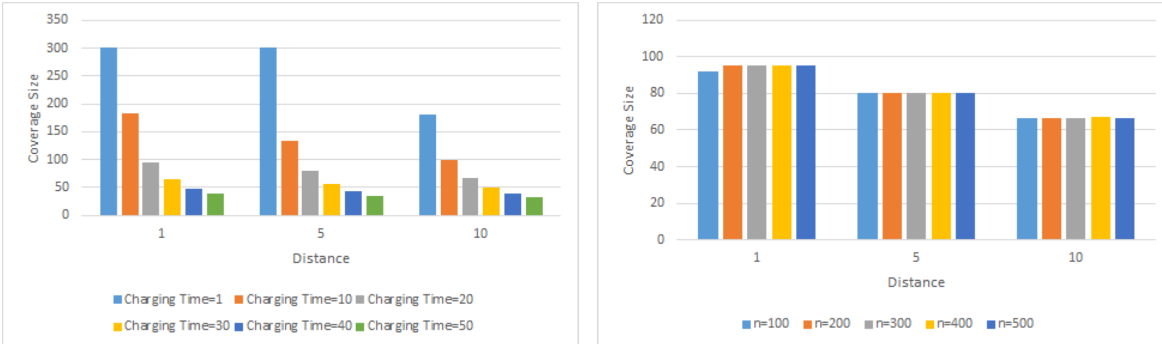

(b) Effect of Travel Time on Coverage Size. $n = 300$ (c) Effect of Travel Time on Coverage Size. $\rho = 20$

Figure 12: Effect of Travel Time on Coverage Size

5.4.5 More detailed analysis.

To better analyses the observed phenomena concerning coverage size, we now focus on some specific scenarios to analyze how the coverage size changes in time.

Starting with arbitrary initial charge levels, the experimental results show that the first few tours of the ring performed by the robot follow a schedule that highly depends on the initial distribution of charges; however, we confirm our conjecture by observing that, for all choices of parameters:

- the charging pattern stabilizes becoming periodic;
- stabilization occurs within two rounds.

In other words, the networks become stable within two rounds. Figure 9, for example, shows coverage size varying ρ confirming that whenever we have $n > \frac{2\Delta}{\rho} + 1$, the average coverage size stabilizes, that coincides with the theoretical value shown in Theorem 5. The same is observed when varying the node capacity Δ.

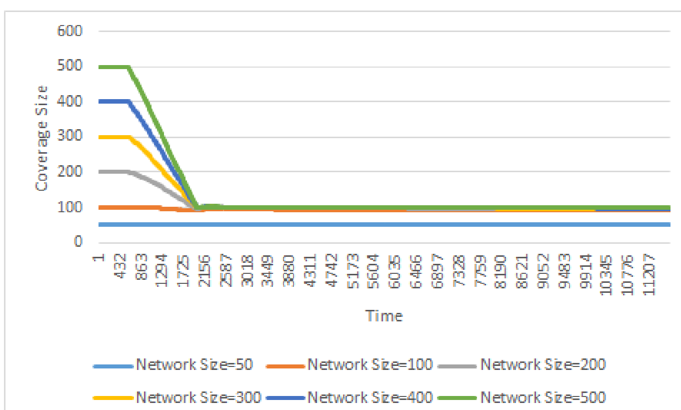

Figure 13: Coverage Size for Δ=2000,d=1,ρ=20.

Figures 13 and 14 also display a clear evidence of stabilization by showing coverage in time for different network sizes under different choices of Δ and ρ. When executing each scenario, computing coverage size and disconnection time, we have also recorded the charging order of the nodes. We have seen that, after at most two rounds, regardless of the size, the network stabilizes. Stability occurs for any value of n, we then can observe that for $n > \frac{2\Delta}{\rho} + 1$ (i.e., $n = 200, 300, 400, 500$ in Figure 13 and all values of n in Figure 14), coverage does not depend on n and coincides with the theoretical value which is shown in the next section.

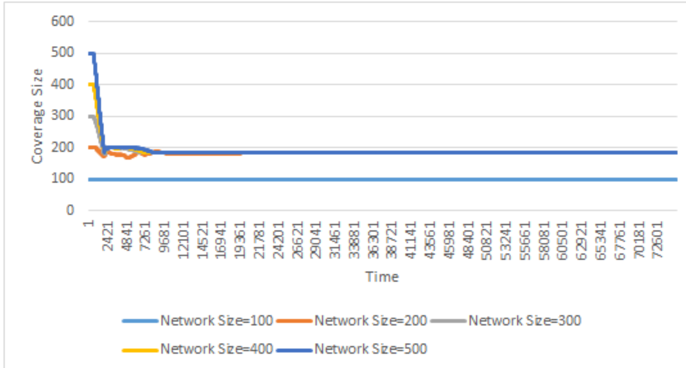

Figure 14: Coverage Size for Δ=2000,d=1,ρ=10.

Finally, we observe that these phenomena hold in all experiments regardless of the initial distribution of charges; that is, the same phenomenon has been observed not only with random distributions, but also with specific ones, such as when the charges are increasing in a clockwise order (see Figure 15) or in counterclockwise order (see Figure 16).

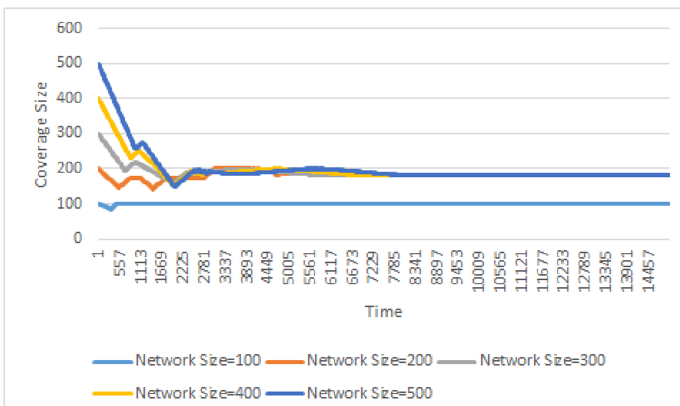

Figure 15: Increasing initial charges

5.5 Comparison with `OPTIMAL`

We now turn to the comparison between `LIC` and `OPTIMAL`. We already established that, when $n > \frac{2\Delta}{\rho} + 1$, the two strategies are equivalent both in terms of coverage size and disconnection time (Theorem 7). Extensive experimental results, varying Δ, ρ, d, and n, confirm the theoretical findings.

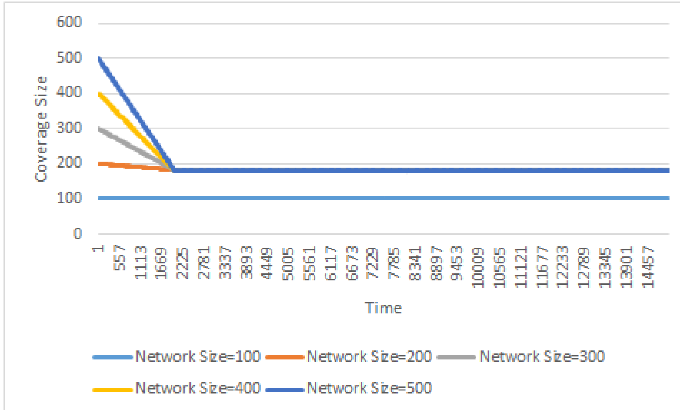

Figure 16: Decreasing initial charges

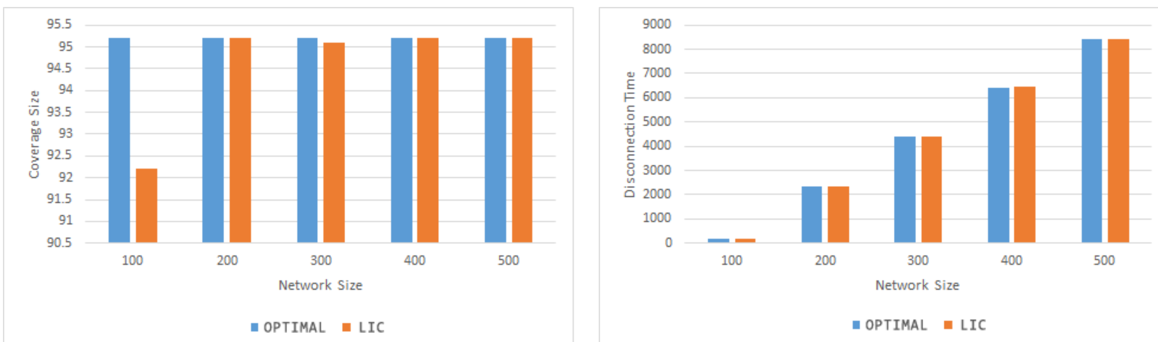

(a) Coverage.

(b) Disconnection time.

Figure 17: `LIC` vs `OPTIMAL` ($\Delta = 2000$, $\rho = 20$, $d = 1$).

For example, Figure 17 shows coverage and disconnection time of the two strategies for different network sizes when $\Delta = 2000$, $\rho = 20$ and $d = 1$. Notice that Theorem 7 does not hold when $n < \frac{2\Delta}{\rho} + 1$; in fact, as shown in Figure 17, `OPTIMAL` has a much better coverage than `LIC` for the case $n = 100$.

6 Concluding Remarks

In this paper, we introduced the notion of effectiveness of energy restoration strategies. We proposed a very simple decentralized battery recharging strategy, which, in spite of its simplicity and of the use of very limited resources, achieves optimal effectiveness in most cases. The technique is based on the on-demand visit of the nodes by a mobile robot in a predefined circular order only when aware of a pending request. The optimality of the strategy is proven for sufficiently large networks ($n > \frac{2\Delta}{\rho} + 1$). It would be interesting to consider also the case of smaller n, where our strategy is not optimal; the detailed analysis of the charging dynamics for that case will be the object of future study.

Our studies, both analytical and experimental, have been carried out in an abstract setting, with several simplifying assumptions. Among them, we assumed that the time necessary for the robot to move from a node to its successor in the cyclic order is uniform. We did run experiments with variable distances between nodes; all these experiments do not show any significant difference with the results obtained in the paper with uniform distances; the theoretical validation is however left for future work.

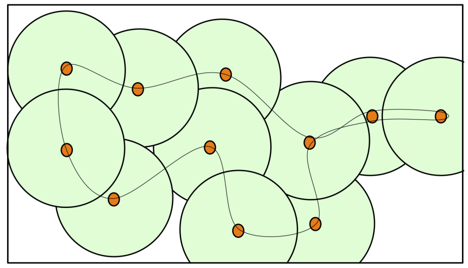

Figure 18: Creating a virtual cycle using DFS.

Our strategy is based on the existence of a cyclic order to be used by the charging robot, where successive nodes in the order are within direct communication range. Should such a cycle not exist in the communication graph (because it is not Hamiltonian) or be difficult to compute, a virtual cycle based on a simple DFS could be used (see Figure 18). The drawback would be that the communication distance between two neighbours in the cycle might be more than one (i.e., multi-hop communication might be required) and the physical distance between them might increase. Alternatives to the cyclic order, to be considered and explored, are less regular ordering structures which decentralized building techniques already exist (e.g., [4]).

Another assumption that would be interesting to lift is the one of constant charging rate ρ, the same for all nodes. The case of variable charging rates, possibly depending on the current battery level, as well as other physical factors (e.g., battery capacity decay) are important open research directions.

Acknowledgment

This research has been supported in part by the Natural Sciences and Engineering Research Council of Canada (N.S.E.R.C.) under the Discovery Grant program.

References

[1] G. Anastasi, M. Conti, M.D. Francesco, A. Passarella, "Energy conservation in wireless sensor networks: A survey". *Ad Hoc Networks*, Vol. 7, no. 3, 2009.

[2] C. M. Angelopoulos, S. Nikoletseas, T. P. Raptis, C. Raptopoulos, and F. Vasilakis, "Improving sensor network performance with wireless energy transfer", *Int. Journal of Ad Hoc and Ubiquitous Computing*, Vol. 20, no. 3, 159-171, 2015.

[3] F. Arwin, K. Samsudin, and A.R. Ramli, "Swarm robots longterm autonomy using moveable charger", in *Proc. International Conference on Future Computer and Communication*, 127–130, 2009.

[4] A. Banerjee and C.-T. King, "Building ring-like overlays on wireless ad hoc and sensor networks". *IEEE Transactions on Parallel and Distributed Systems*, Vol. 20, No. 11, 1555 –1566, 2009.

[5] H. Dai, X. Wu, L. Xu, G. Chen, and S. Lin, "Using minimum mobile chargers to keep large-scale wireless rechargeable sensor networks running forever", in *Proc. 22nd International Conference on Computer Communications and Networks (ICCCN)*, 1–7, 2013.

[6] A. Drenner and N. Papanikolopoulos, "Docking station relocation for maximizing longevity of distributed robotic teams", in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*. 2436–2441, 2006.

[7] L. Fu, P. Cheng, Y. Gu, J. Chen, and T. He, "Optimal charging in wireless rechargeable sensor networks", *IEEE Transactions on Vehicular Technology*, Vol. 65, No. 1, 278–291, 2016.

[8] S. He, J. Chen, F. Jiang, D. K. Yau, G. Xing, and Y. Sun, "Energy provisioning in wireless rechargeable sensor networks," *IEEE Transactions on Mobile Computing*, Vol. 12, No. 10, 1931–1942, 2013.

[9] L. He, Y. Gu, J. Pan, and T. Zhu, "On-demand charging in wireless sensor networks: Theories and applications" in *Proc. 10th IEEE International Conference on Mobile Ad-Hoc and Sensor Systems (MASS)*, 28–36, 2013.

[10] L. Jiang, X. Wu, G. Chen, and Y. Li, "Effective on-demand mobile charger scheduling for maximizing coverage in wireless rechargeable sensor networks", *Mobile Network Applications*, Vol. 19, 543–551, 2014.

[11] L. Khelladia, D. Djenouria, M. Rossib, N. Badach, "Efficient on-demand multi-node charging techniques for wireless sensor networks", *Computer Communications*, Vol. 101, 44-56, 2017.

[12] S. Kumar, T.H. Lai, and A. Arora, "Barrier coverage with wireless sensors", *11th International Conference on Mobile Computing and Networking (MobiCom)*, pp. 284–298, 2005.

[13] A. Kurs, A. Karalis, R. Moffatt, J. Joannopoulos, P. Fisher, M. Soljacic, "Wireless power transfer via strongly coupled magnetic resonances", *Science*, Vol. 317, 83–86, 2007.

[14] A. LaMarca, W. Brunnete, D. Koizumi, M. Lease, S.B. Sigurdsson, K. Sikorski, D. Fox, and G. Borriello, "Making sensor networks practical with robots", in *Proc. International Conference on Pervasive Computing*, 152–166, 2002.

[15] X. Li, R. Falcon, A. Nayak, and I. Stojmenovic, "Servicing wireless sensor networks by mobile robots", *IEEE Communications Magazine*, Vol. 50, No. 7, 147-154, 2012.

[16] W. Liang, W. Xu, X. Ren, X. Jia, and X. Lin, "Maintaining large-scale rechargeable sensor networks perpetually via multiple mobile charging vehicles", *ACM Trans. on Sensor Networks*, Vol. 12, No. 2, 14:1-14:26, 2016.

[17] L. Liu, B. Hu, and L. Li,"Energy conservation algorithms for maintaining coverage and connectivity in wireless sensor networks. *IET communications*, Vol. 4, No.7, 786–800, 2010.

[18] S. Luke, C. Cioffi-Revilla, L. Panait, K. Sullivan, and G. Balan, "Mason: A multiagent simulation environment", *Simulation*, Vol. 81, No. 7, 517–527, 2005.

[19] A. Madhja, S. Nikoletseas, T.P. Raptis, "Distributed wireless power transfer in sensor networks with multiple mobile chargers", *Computer Networks*, Vol. 80, 89–108, 2015.

[20] F. Michaud and E. Robichaud, "Sharing charging stations for long-term activity of autonomous robots", In *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2746–2751, 2002.

[21] L. Militano, M. Erdelj, A. Molinaro, N. Mitton, and A. Iera, "Recharging versus replacing sensor nodes using mobile robots for network maintenance", *Telecommunication Systems*, Vol. 63, No. 4, 625–642, 2016.

[22] N. Mitton and D. Simplot-Ryl (Eds.), *Wireless Sensor and Robot Networks*, World Scientific, 2014.

[23] Y. Mei, C. Xian, S. Das, Y. Hu, and Y. Lu, "Sensor replacement using mobile robots", *Computer Comm.* Vol. 30, No. 13, 2615–2526, 2007.

[24] E. Omar, *Energy Restoration of Sensor Networks by Mobile Robots*, PhD thesis, School of Electrical Engineering and Computer Science, University of Ottawa, 2019.

[25] E. Omar, P. Flocchini, and N. Santoro, "Energy restoration in a linear sensor network", *Proc. 15th IEEE International Conference on Computer Systems and Applications*, 1–8, 2018.

[26] Y. Peng, Z. Li, W. Zhang, and D. Qiao, "Prolonging sensor network lifetime through wireless charging", in *Proc. 31st IEEE Real-time Systems Symposium (RTSS)*, 129–139, 2010.

[27] M. Rahimi, H. Shah, G. Sukhatme, J. Heidemann, and D. Estrin, "Studying the feasibility of energy harvesting in a mobile sensor network", in *Proc. IEEE International Conference on Robotics and Automation*. 19–24, 2003.

[28] N. Santoro and E. Velazquez, *Energy Restoration in Mobile Sensor Networks*, Chapter 6 of [22], World Scientific, 2014.

[29] M. Sharifi, S. Sedighian, and M. Kamali, "Recharging sensor nodes using implicit actor coordination in wireless sensor actor networks", *Wireless Sensor Network*, vol. 2, 123–128, 2010.

[30] Y. Shi, L. Xie, Y. T. Hou, and H. D. Sherali, "On renewable sensor networks with wireless energy transfer", in *Proc. 30th IEEE International Conference on Computer Communications (INFOCOM)*, 1350–1358, 2011.

[31] S. Sudevalayam and P. Kulkarni, "Energy harvesting sensor nodes: Survey and implications". *Communications Surveys and Tutorials*, Vol. 13, 1–19, 2011.

[32] B. Tong, G. Wang, W. Zhang, and C. Wang. "Node reclamation and replacement for long-lived sensor networks", *IEEE Transactions on Parallel and Distributed Systems*, Vol. 22, No. 9, 1550–1563, 2011.

[33] E. Velazquez, and N. Santoro, "Mobility-based strategies for energy restoration in wireless sensor networks", *6th International Conference on Mobile Ad-hoc and Sensor Networks (MSN)*, 161–168, 2010.

[34] C. Wang, J. Li, F. Ye, and Y. Yang, "Multi-vehicle coordination for wireless energy replenishment in sensor networks", in *Proc. 27th IEEE International Symposium on Parallel and Distr. Processing (IPDPS)*, 1101–1111, 2013.

[35] W. Xu, W. Liang, X. Ren, X. Lin, "On-demand energy replenishment for sensor networks via wireless energy transfer", in *Proc. IEEE International Symposium on Personal Indoor Mobile Radio Comm. (PIMRC)*, 1269-1273, 2014.

[36] L. Xie, Y. Shi, Y.T. Hou, and H.D. Sherali, "Making sensor networks immortal: an energy-renewal approach with wireless power transfer", *IEEE/ACM Transactions on Networks*, Vol. 20, No. 6, 1748–1761, 2012.

[37] S. Zhang, J. Wu, and S. Lu, "Collaborative mobile charging for sensor networks," *9th IEEE International Conference on Mobile Adhoc and Sensor Systems (MASS)*, 84–92, 2012.