

Using Mobile Robots to Harvest Data from Sensor Fields

Onur Tekdas[‡] Jong Hyun Lim[†] Andreas Terzis[†] Volkan Isler^{‡*}
University of Minnesota[‡] Johns Hopkins University[†]

November 29, 2008

Abstract

We explore synergies among mobile robots and wireless sensor networks in environmental monitoring through a system in which robotic data mules collect measurements gathered by sensing nodes. A proof of concept implementation demonstrates that this approach significantly increases the system's lifetime by conserving energy that the sensing nodes would otherwise use for communication.

1 Introduction

The recent emergence of small embedded sensing devices (also known as *motes*) will fundamentally change the way we approach many scientific problems, some of them totally intractable in the past. One such family of scientific problems is collectively known as *environmental monitoring*, in which the scientists must collect environmental data (e.g., temperature, humidity, etc.) over long temporal and spatial scales [6]. Such experiments today are infeasible because they require significant manual labor and perturb the environment under observation.

With sensor network technology, scientists can deploy sensing motes close to the phenomenon they want to observe and collect measurements from the motes' onboard sensors. Because the motes are physically small, battery operated, and include a wireless radio, deploying such a *wireless sensor network* perturbs the environment minimally and reduces the installation and maintenance costs. Furthermore, the inexpensive nature of these devices enables scientists to place a high resolution mote grid in the field, and get frequent measurements, providing an extremely rich data set about the dependencies and subtle differences among many correlated parameters.

Nonetheless, multiple challenges remain before the promise of autonomous environmental monitoring can be fulfilled. Limited energy is one of the key challenges, considering the motes are most often battery operated. The mote's wireless radio is its largest energy consumer, other than possibly its on board sensors. This observation means that the motes can spend the majority of their energy stores transmitting collected data, especially because they may need to *forward* data for other motes in the network. Moreover, considering that wireless links are notoriously fickle and scientists want all the collected data, retransmissions are likely, further increasing the energy cost.

In this paper we present an alternative to an end-to-end wireless network that forwards the motes' measurements to a back-end database. Instead, we propose to use autonomous robots as *data mules*. These mules visit locations within the communication distance of each of the static motes, download their measurements and return to a remote base station to offload the collected data. The key benefit of this approach is that motes can conserve energy that they would otherwise use to forward data, thereby prolonging the network's lifetime. However, in order to successfully deploy such a hybrid robot/sensor network, new challenges (such as planning robots' trajectories) must be overcome.

The remainder of this paper summarizes background information on environmental sensor networks, exemplified by a network we have deployed, and presents the challenges associated with collecting data from remotely deployed networks. Next, we describe how robots can assume the role of data mules and the technical challenges associated with this task. Finally, we present a proof of concept system that we developed to explore the trade-offs in using robots as data mules.

*Isler and Tekdas were at Rensselaer Polytechnic Institute when this work was performed. Their work is supported in part by NSF IIS-0745537, NSF CCF-0634823 and NSF CNS-0707939

2 Related work

The idea of exploiting mobility for data collection and routing has received some recent attention. In most of the existing work, data mules are treated as uncontrolled agents. In this section, we present a brief overview of related work focusing on controllable agents. Yuan et al. formulate the problem of collecting sensor data using a single robot as a special instance of the Traveling Salesperson Problem with Neighborhoods (TSPN), where each neighborhood is a disc whose center corresponds to a node and radius is determined by the communication range [11]. We present algorithms for multiple robots and focus on systems challenges associated with energy considerations. Tirta et al. presented algorithms to schedule visits of a mobile agent to collect data from cluster heads [10]. The authors present heuristics which focus on data latency and data aggregation rate of clusters. In our work, we focus on energy efficiency issues. Another aspect that distinguishes our work from the previous body of work is that we present a proof of concept implementation using robots and motes.

One of the most related results is presented in [3]. In that work, Dunbabin et al. present an underwater data muling system. We focus on a land-based system. The corresponding system designs reflect the challenges associated with the two scenarios. In the underwater scenario, sensors and underwater vehicle communicate through optical communication, which requires a close proximity as well as good view-angle to start the communication. Moreover, since GPS localization is not available under the water, the vehicle has to navigate under high localization uncertainty. This makes designing global routing algorithms challenging. The authors propose a solution where the vehicle performs a spiral movement to find the sensors. This strategy is not efficient for our scenario, in which the sensor locations are known and the robot can localize itself. The authors present a real implementation with three sensors and a single robot. We present global routing algorithms for larger sensor networks which utilize multiple robots (up to three).

3 Background

A wireless sensor network (WSN) is a collection of small embedded computing devices (motes) deployed over an area of interest, communicating through wireless radios. These motes consist of a microcontroller, data storage devices, sensors, analog-to-digital converters (ADCs), a data transceiver, and an energy source (e.g., AA batteries). Existing motes use an 8 or 16-bit microcontroller with tens of KBytes of RAM, hundreds of KBytes of ROM for program storage and external storage in the form of FLASH memory [7]. These devices operate at one milliwatt while running at about 10 MHz. Most of the circuits can be powered off, so the standby power can be about one microwatt. If such a device is active 1 percent of the time, its average power consumption is just a few microwatts enabling long term operation even with two AA batteries. Mote radios transmit at rates between 10-250 Kbps, consume about 20 milliwatts when in transmit, receive, or idle listening mode, while their range typically is measured in tens of meters.

One of the most popular wireless sensor network applications is environmental monitoring [6, 8]. In its basic form, environmental monitoring involves reliably gathering the measurements from each of the network's motes. Furthermore, in order to observe long term spatial and temporal trends, these networks need to be deployed for years and cover large geographic areas. These requirements mean that such networks must have duty cycles of 1% or lower and be deployed sparsely over the monitored area.

One such network is a soil monitoring network we deployed at the Jug Bay wetlands sanctuary. This sanctuary is located along the Patuxent river in Maryland and serves as the habitat for a variety of turtle species, including the Eastern Box turtle. These turtles are of scientific interest because their sex is not determined by sex genes, but by the eggs' incubation temperature. It has been shown in the lab that a difference of two degrees centigrade is enough to produce male instead of female offsprings. On the other hand, the in vivo conditions of box turtle nests have not been observed in the wild. For this reason we deployed a mote network that monitors soil temperature and soil humidity at turtle nests. Because we have no control of where turtles decide to lay their eggs, our motes are spread over a very large geographic area (approximately 25 square kilometers).

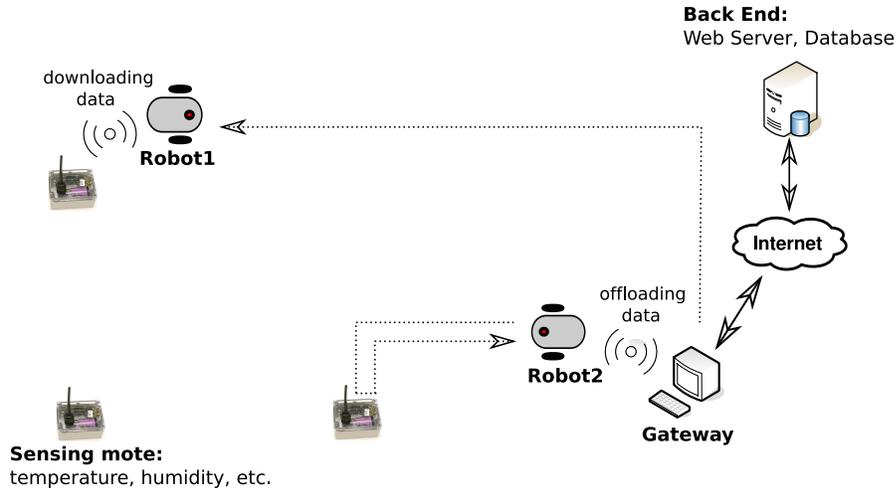


Figure 1: Overall system architecture, including a number of sensing motes, multiple robots acting as data mules, and a gateway to which robots offload collected data and receive further commands. Each robot communicates with the sensing nodes and the gateway through a locally connected mote. All collected data are eventually stored in a back-end database for further processing and visualization.

3.1 Data retrieval

It is uneconomical to deploy an end-to-end environmental monitoring WSN over a large geographic area. The reason is that, depending on the environment, the range of individual wireless links is between 10 and 50 meters. Consequently, a large number of *relay points* is necessary for multi-hop data collection. Even if deploying such a network were economically viable, reliably delivering data over it would consume considerable energy. This is because a single packet needs to be transmitted multiple times, by each mote on the routing path to the gateway. Furthermore, because wireless links are lossy, the probability of loss (and thus retransmission) increases as the length of the routing path increases. Finally, in addition to the energy consumed to carry sensor measurements, the motes need to exchange control traffic to maintain the end-to-end routing paths used to forward data traffic, further depleting their limited energy stores.

On the other hand, mobile robots can navigate the terrain, move physically close to each of the network's motes and download their collected measurements over a direct wireless link. Once the robots collect the desired data from the network they return to the gateway, where they offload them. We term the robots that collect and ferry data as *data mules*. Figure 1 graphically depicts the proposed system architecture. This architecture offers multiple advantages. First, data mules can move close enough to a mote to ensure that the quality of the wireless link will be high. Furthermore, because the robot moves adjacent to the mote, the mote can reduce its radio's transmission power, further reducing its energy consumption. Finally, recharging the robots' batteries is a simpler operation than replacing motes' batteries. At the same time, using data mules introduces its own set of challenges which we outline in the section that follows.

4 Mobile robots as data mules

4.1 Node wakeup and data extraction

As previously mentioned, sensing motes keep their radios mostly off to conserve energy. Then to check for any pending requests, each sensing mote transmits a beacon and waits for an acknowledgment. If no acknowledgment arrives the mote turns its radio off, otherwise it remains active waiting for data download requests. This *Low Power Probing* (LPP) mechanism was initially proposed in [5]. By controlling the beacon frequency the node can control its duty cycle. For example, selecting a beaoning interval of 20 seconds,

leads to a $\sim 0.1\%$ duty cycle, as a beacon-acknowledgement cycle lasts 20 msec.

Based on this description, the robot moves close to the sensing mote, waits for a beacon message, acknowledges it, and then issues a command to the mote to download (some of) its data. We can further reduce the mote’s energy consumption if the schedule of the robot’s visits is known. In this case, the robot informs the mote when it will return. The mote then broadcasts its beacons only shortly before until shortly after the robot is expected to return. Doing so reduces the energy consumed to send the beacons at the cost of reducing the system’s flexibility (the robot can arrive only during pre-described times). If the mote loses the robot’s schedule (e.g., after an unexpected reboot operation), it can fall back to its periodic probing schedule.

An alternative wakeup strategy is to use a lower power sensor to wake up the node. For example one could use the mote’s light sensors or even a Reed switch (i.e. an electrical switch operated by an applied magnetic field) to wake up the mote. In this case, the robot could shine a light or carry a magnet that would activate the sensor which will in turn activate the sensing mote’s radio for the subsequent download.

Once the mote is awake, the robot transmits a request containing the range of collected data it needs to download from the mote. The actual data download process uses the standard NACK-based Automatic Repeat reQuest (ARQ) protocol.

4.2 Location Information

Throughout this work, we assume that accurate location information, for both the robots and motes, is available. For most applications, it is reasonable to assume that robots can self-localize at all times. Specifically, robots can use GPS when outdoors, whereas in indoor environments, in which visual features abound, vision-based localization techniques can be utilized. Likewise, one can assume that robots have complete and accurate knowledge of the motes’ locations, as these locations can be recorded during the deployment process. Nonetheless, in the paragraphs that follow we briefly discuss how this assumption can be relaxed.

The accuracy of off-the-shelf GPS devices is between 3-5 m, while differential GPS devices reduce the error to ~ 20 cm. A communication link between a mote and a data mule in the presence of location uncertainties can be established by using a longer communication range. For example, suppose we are using a GPS device with ~ 3 m localization error. Further, let us suppose that the locations of the motes were recorded using this device, and that the robot uses an identical GPS for navigation. In this case, we can set the communication range of the motes to ~ 6 m. This guarantees that even if a mote and a robot are at the opposite sides of their uncertainty regions, they can still establish a connection. On the other hand, increasing the communication range does not guarantee error-free communications due to *multipath interference*. We are currently investigating efficient algorithms to overcome this issue by using the robot’s mobility to find locations within the mote’s communication range that ensure low packet loss.

Moreover, if the locations of motes are not known, an interesting “network discovery problem” arises in which a discovery phase is implemented. During this phase, motes periodically transmit beacons while the robots cover the environment. When a robot listens a beacon, it learns an approximate estimate of the mote’s location (or range) which it can propagate to the other cooperating robots. Similar spatial coverage problems have recently received significant attention in the robotics literature. However, the network discovery problem mentioned above has a temporal aspect as the motes’ messages arrive at discrete time-steps. This calls for design of new spatio-temporal coverage algorithms which guarantee that all motes are discovered as quickly as possible.

4.3 Robot Routing

The task of collecting sensor data with mobile robots introduces a number of challenging path planning problems, related to fundamental vehicle routing problems. In the problem’s most basic form, we have a single robot which must visit mote locations as quickly as possible. This is the well-known *Traveling Salesperson Problem (TSP)* which asks for the shortest route for a salesperson (the robot) to visit a set of cities (the motes). In the open field, data aggregation with robots corresponds to a special case of TSP, called the *Euclidean TSP* where the edge cost between two nodes is equal to their Euclidean distance. Efficient approximation algorithms for Euclidean TSP exist in the literature. [1].

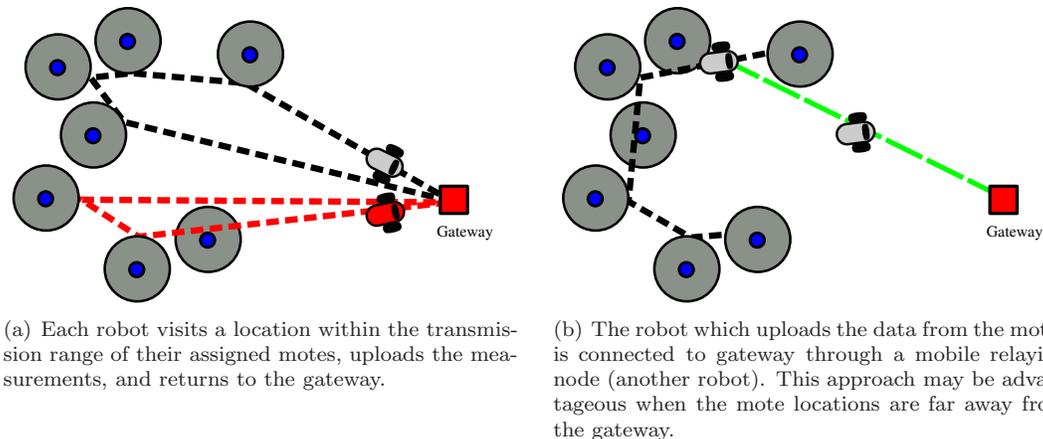


Figure 2: Two alternative download strategies. The small blue circles correspond to motes, while the larger gray circles correspond to the transmission region of a mote. In reality the transmission range is of course not circular, but robots can move close enough to the motes to avoid any RF propagation irregularities.

For one data mule, a TSP tour gives an optimal strategy to collect the motes' data. However, one can increase the aggregation frequency and reduce the distance the mule travels, by employing multiple data mules and splitting the tour over the number of data mules. Computing tours for k agents is known as the k -TSP problem [4]. An optimal k -TSP solution finds k tours, where the largest of k tours is minimized. Since the robots in our system will establish wireless connections with the sensors, a robot does not need travel all the way to a sensor's location to download its data. This capability can be modeled as a *TSP with Neighborhoods (TSPN)* problem. In TSPN, each customer is willing to travel a certain distance to meet with the salesperson. This corresponds to finding the shortest route visiting a given set of disks on the plane (see Figure 2(a).)

So far, we focused on designs in which motes are assigned to robots and the robots are used primarily for data collection. In some scenarios, it is possible to improve the system performance (e.g., latency) and limit the robots' energy consumption by enabling *multi-robot collaboration*. One such scenario is when the gateway location is significantly further away from the sensors. In this case, instead of having robots visit the gateway periodically, some robots can act as *robotic routers* [9] and relay messages to the gateway. Doing so will significantly reduce the time spent traveling (see Figure 2(b).)

While approximation algorithms exist for the standard TSP variants, additional requirements must be introduced to model the temporal aspects of the data gathering application. For example, the size of a mote's memory introduces an upper-bound on the length of the time between robots' consecutive visits to the mote. Furthermore, synchronizing robot's arrival and mote's wake-up time, when motes do not constantly transmit beacons, requires introducing time windows within which the robot must arrive at the sensors. Recent results on variations of the data gathering problem with time constraints can be found in [2].

5 Proof of Concept

Next, we present a proof of concept system we developed, as well as results from a small experiment that proves the feasibility of using mobile robots as data mules.

Our proof of concept consists of a network of sensing motes and a number of mobile robots. Each sensing mote is a Tmote Sky from Moteiv [7]. Moreover, each of the Acroname Garcia robots has a Stargate PDA-class single board computer, that controls the robots' motors and sensors (see Figure 3). The Stargate uses a Tmote Sky motes, directly connected to its USB interface, to communicate with the sensing nodes. The communication between the mote and the robot's computer follows a simple command-and-response protocol that we developed. Once the robot approaches a sensing node, it waits for a beacon message to verify that the node is still alive. After it receives the beacon, the robot initiates the data download process. If on the other hand, a beacon is not received within a predetermined amount of time, the robot proceeds



Figure 3: An Acroname Garcia robot visiting a sensing mote. Note the single board computer on the robot under the clear lid and the mote connected to the robot’s computer.

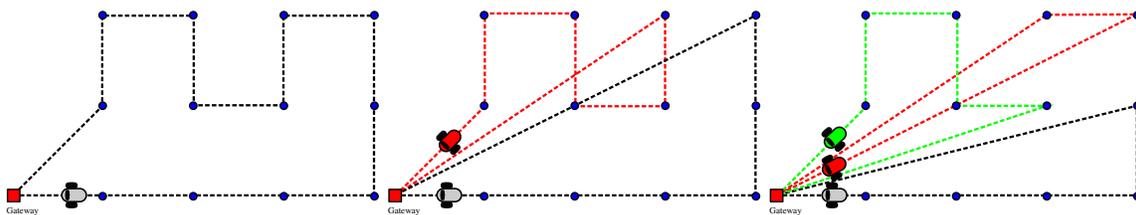


Figure 4: From left to right: Optimal TSP tours for one, two, and three data mules. The red square represents the gateway and blue circles represent mote locations.

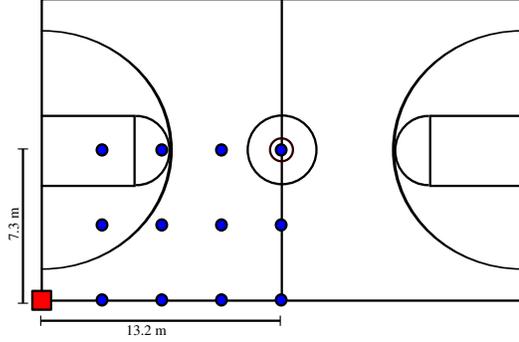
to the next sensing mote on its path.

We assume that the motes’ transmission power is set to its lowest level to reduce energy consumption. Hence, while planning the mules’ trajectories, we set the visiting locations to be the motes’ exact locations. The trajectories of the multi-robot system are pre-calculated as k -TSP tours. Figure 4 illustrates the resulting k -TSP tours for $k = \{1, 2, 3\}$ on a regular grid where twelve motes are deployed. For this small TSP instance, we find the optimal TSP tour by enumeration, which is feasible for this rather small network. For larger problems, efficient TSP solvers can be used. To compute the 2-TSP and 3-TSP tours, we split the optimal TSP tour into smaller tours using the k -SPLITOUR algorithm [4].

We illustrate the practical feasibility of data gathering using data mules with a small-scale experiment. For this experiment, we deployed twelve motes and a gateway on a virtual grid on a basketball field at the Rensselaer Polytechnic Institute (see Figure 5(a)). The robot system we developed for the experiment follows a step by step procedure. In each step, a robot takes a turn or takes a forward step with a length of at most 7 m. We limit the length of each step to limit the robot’s deviation from its designed path. Such deviations occur due to the inherent accuracy limits of the robot’s motors. Because the goal of this experiment is not to evaluate the accuracy of indoor localization techniques and autonomous repositioning, we manually re-position the robots to correct the position if they veer off the desired path. In total, we conducted three experiments with one to three data mules. Figure 4 shows the robots’ trajectories for each of the three experiments and Figure 5(b) shows a snapshot from the experiment with two data mules.

5.1 Evaluation

As described in Section 3.1, one key reason for using data mules is to reduce the sensing motes’ energy consumption and thus prolong the network’s lifetime, while guaranteeing reliable data collection. We quantify this benefit by comparing the energy consumed to send a single packet from each mote in two settings. In



(a) Map of the deployment area. The deployment consists of twelve motes (blue circles) and a gateway (red square), placed on a grid that spans a quarter of a basketball court.



(b) A snapshot from the experiment with two mules. One mule (white robot) is following its designed path while the other mule (red robot) is downloading data from the assigned mote.

Figure 5: The map of our deployment (left) and a snapshot from the experiment with two mules (right).

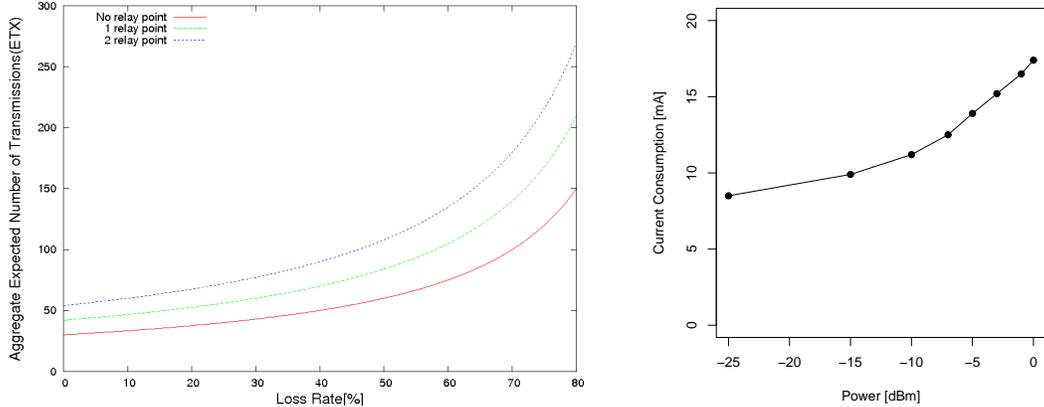
the first setting, an end-to-end, multi-hop wireless network system is used to collect the data. In the second setting, sensing motes directly offload their packets to data mules as proposed in this paper. We simulate this comparison based on a grid of twelve motes whose pattern mirrors the 3×4 grid that we deployed in the basketball court. The deployment is shown in Figure 5(a). On top of this, we make the following assumptions. First, we assume that each mote can communicate only with its direct neighbors on the grid and therefore the longest end-to-end network path spans four wireless links. Second, each packet can be lost with probability p , which is uniform across all of the network’s links. Finally, we use a hop-by-hop retransmission scheme in which case each node will retransmit the packet if it does not receive an acknowledgment from its upstream neighbor.

Given these assumptions, we calculate the sum of the expected number of transmissions (ETX) necessary to have each of the network’s twelve motes successfully deliver a single packet to the gateway. Based on the end-to-end routing tree overlayed on the mote grid, the aggregate ETX is equal to 30 when $p = 0$ (i.e., no packet loss). On the other hand, when data mules are used, the aggregate ETX is equal to 12 since each mote can directly offload its packet to a mule over a short and reliable wireless link. While the benefits of using data mules are clear even with error-free radio links, as Figure 6(a) suggests, the disparity rapidly increases as p increases. Furthermore, this estimate assumes that the gateway can directly communicate with some of the network’s motes. In reality, however, the gateway may not be in sensing motes’ communication range in which case one or more *relay points* are necessary. Adding relay points, however, increases the length of the end-to-end path and thus further increases the ETX.

The second benefit of using data mules is that motes can decrease their transmission power and thus save energy. Figure 6(b) supports this claim, showing that if a mote reduces its transmission power level from its default level of 1 mW (0 dBm) to 0.003 mW (-25 dBm), the energy consumption is reduced drastically. Furthermore, because the mules can reliably approach sensing motes to a close distance, reducing the transmission power does not increase the probability of loss.

In addition to estimating the energy the sensing motes consume, we also consider the robot’s energy consumption. Since robot batteries are also limited, yet rechargeable energy sources, we need to consider how long the batteries last when we dispatch the robots to collect data. If the network is too big and therefore the tour is too long, the energy stored in a single robot battery may not be sufficient to complete a single tour. In that case, we can split the the single tour until we get feasible tours that each robot can complete at least once, before recharging its battery. Even if a single robot can tour the whole network, we need to be able to predict how many tours it can complete before it needs to recharge its battery. Hence, accurately estimating the robot’s energy consumption is a prerequisite for the correct operation of the data muling system.

The robots we use are powered by standard, six-cell 7.6 V 3,000 mAh NiMH battery packs. The current drawn from the battery varies depending on the robot’s activity. We classify these activities in three states: *stand-by*, *moving*, and *downloading/uploading data*. We experimentally measured the minimum and



(a) Total number of expected transmissions for a 3×4 grid as a function of packet loss probability p . Each mote delivers a single packet to the gateway. The different lines correspond to grids that require zero to two relay points to connect to the network's gateway.

(b) Current consumed by the mote's radio (TI Chipcon CC2420) as a function of transmission power.

Figure 6: Using data mules reduces the sensing mote's energy consumption in two ways: by reducing the total number of transmissions necessary to deliver measurements to the database and by reducing the sensing motes' transmission power.

maximum current values for each of these three states. The measured values are: 560 ~ 620 mA (stand-by), 630 ~ 940 mA (moving), and 590 ~ 750 mA (download/upload).

During our experiments, robots had an average speed of 0.2 m/sec. Downloading 11.4 KBytes of data from each mote took one minute in average, which consists of 20 seconds beaconing and actual data download time. In the first experiment, the trajectory length for one data mule was 46.2 m. The total time used by the robot to follow the trajectory, download data from the motes, and upload to the gateway was 16:51 minutes. In the second experiment, we used two robots where the robots' trajectory lengths were 35.5 m and 34.7 m. The total times were 9:56 min and 9:57 min, respectively. Finally, in the third experiment, we used three robots with trajectory lengths of 30.5 m, 30.6 m, and 29.3 m. The total times in that case were 8:33 min, 5:33 min and 8:26 min, respectively.

While calculating the energy consumption, we consider the maximum measured current values. According to this assumption, the robot consumed 1.69 Wh in the first experiment. In the second experiment, both robots consumed 1.02 Wh, while in the third experiment, robots consumed 0.87 Wh, 0.59 Wh and 0.86 Wh, respectively. One can then use these estimates and the battery's capacity (22.8 Wh), to determine how many tours the robots can complete before they need to recharge their batteries.

6 Discussion

Wireless sensor network technology has the potential to enable major breakthroughs in natural sciences by giving scientists the capability to collect high fidelity data over large geographic regions and extended periods of time. We argue that mobile robots can help sensor networks achieve their full potential. In this paper, we explored one such synergy between robots and a static sensor field, in which robots act as mules that collect the measurements acquired by the sensors. We used a proof-of-concept system to show that this approach is feasible and yields important savings in energy costs, prolonging the network's lifetime.

A number of improvements are necessary before our system can be fully autonomous. Specifically, we are currently extending the localization capabilities of our robots by implementing vision-based localization techniques. In motion planning, we plan to improve our system in two ways: (i) Currently, our robots visit the exact location of the motes. We are planning to utilize the communication range of the motes by implementing TSPN-based approaches in route planning. (ii) In our current approach, each mote is assigned to a single robot for data download. We are also planning to investigate the utility of allowing multiple

robots to visit a single mote (at different times), as well as allowing robots to hand off collected data. In the long run, we expect to perform outdoor experiments with real applications. This requires more sophisticated robots capable of navigating outdoor terrains.

References

- [1] S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM*, 45(5):753–782, 1998.
- [2] N. Bisnik, A. Abouzeid, and V. Isler. Stochastic event capture using mobile sensors subject to a quality metric. *IEEE Tran. on Robotics*, 23(4):676 – 692, 2007.
- [3] M. Dunbabin, P. Corke, I. Vasilescu, and D. Rus. Data muling over underwater wireless sensor networks using an autonomous underwater vehicle. *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2091–2098, 15-19, 2006.
- [4] G. N. Frederickson, M. S. Hecht, and C. E. Kim. Approximation algorithms for some routing problems. *SIAM Journal on Computing*, 7(2):178–193, 1978.
- [5] Răzvan Musăloiu-E., Chien-Jan Mike Liang, and Andreas Terzis. Koala: Ultra-Low Power Data Retrieval in Wireless Sensor Networks. In *Proceedings of the Seventh International Conference on Information Processing in Sensor Networks (IPSN)*, April 2008.
- [6] Răzvan Musăloiu-E., Andreas Terzis, Katalin Szlavecz, Alex Szalay, Joshua Cogan, and Jim Gray. Life Under your Feet: A Wireless Sensor Network for Soil Ecology. In *Proceedings of the 3rd EmNets Workshop*, May 2006.
- [7] Joseph Polastre, Robert Szewczyk, and David Culler. Telos: Enabling Ultra-Low Power Wireless Research. In *Proceedings of the Fourth International Conference on Information Processing in Sensor Networks: Special track on Platform Tools and Design Methods for Network Embedded Sensors (IPSN/SPOTS)*, April 2005.
- [8] L. Selavo, A. Wood, Q. Cao, T. Sookoor, H. Liu, A. Srinivasan, Y. Wu, W. Kang, J. Stankovic, D. Young, and J. Porter. LUSTER: Wireless Sensor Network for Environmental Research. In *Proceedings of the 5th ACM Sensys Conference*, November 2007.
- [9] O. Tekdas and V. Isler. Robotic routers. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 1513 – 1518, May 2008.
- [10] Y. Tirta, Zhiyuan Li, Yung-Hsiang Lu, and S. Bagchi. Efficient collection of sensor data in remote fields using mobile collectors. *International Conference on Computer Communications and Networks (ICCCN)*, pages 515–519, Oct 2004.
- [11] B. Yuan, M. Orlowska, and S. Sadiq. On the optimal robot routing problem in wireless sensor networks. *IEEE Trans. on Knowl. and Data Eng.*, 19(9):1252–1261, 2007.