Algorithms for Collecting Data from Cooperating Sensor Motes using Unmanned Vehicles

Kaarthik Sundar[1], P.B. Sujit[2,4], Sivakumar Rathinam[1], Daniel E. Lucani[3,4], Joao Sousa[4]

***Abstract*— This article addresses a fundamental resource allocation problem that arises in monitoring applications: given the locations of the motes, the amount of data that needs to be transferred from each mote to the base station, and an Unmanned Vehicle (UV), find (i) a communication network among the motes, (ii) a subset of motes, referred to as cluster-heads, that act as relays between the motes and the UV, and (iii) a path for the UV such that each mote uses the communication network to transmit its data to one of the cluster-heads, each of the cluster-heads is visited by the UV, and the sum of communication costs involved in transmitting the data from the motes to the cluster-heads and the travel costs of the UV is a minimum. This problem is a generalization of a single Traveling Salesman Problem (TSP) and is NP-Hard. This article presents a rounding algorithm and heuristics to solve the problem. Computational results show that the rounding algorithm performed the best for the tested instances with up to 50 motes and produce solutions that are on an average within 5% of the optimum time relatively fast.**

I. Introduction

Recent advances in small sensing devices, wireless networking and Unmanned Vehicles (UVs) have provided a new way of collecting useful information in environmental monitoring applications [1], [2], [3], [4], [5]. These applications frequently require collecting data such as soil moisture content, temperature and humidity over large swathes of land. Small sensing devices (also referred to as *motes*) can be easily deployed in these applications to collect and transmit the relevant data without disturbing the environment. A mote typically consists of a micro-controller, a wireless radio, data storage devices, sensors and batteries. It can communicate its sensed information either to the UVs or to its neighboring motes. The UVs can then deliver the sensed data from the motes to the base station for further processing.

This approach of using both stationary motes and UVs to collect data is advantageous for several reasons. Firstly, direct communication from the sensed sites to the base station may require a high-power transmitter and may not be suitable for environments with obstructions or non-line-of-sight communications. Simulations/experiments [6], [7] have shown that this type of transmission is also inefficient in terms of energy consumption. Secondly, even if the sensors communicate with the base station through a series of relays (a relay is any device that can receive data from the motes and transmit it; a mote can also perform the role of a relay), power consumption may be high as environmental applications require sensing and communicating over thousands of hectares of land. Relays may also have to only depend on battery power for communication as they may be stationed in areas where direct power from the grid is not available. A UV can travel to the monitoring sites and download the sensed data from the motes, thus reducing the power expended by the motes in relaying large amounts of data. This process can directly help in increasing the life span of the motes. By also using UVs to collect data, the motes are not required to form a connected network and can be spatially distributed depending on the constraints of the application.

This article addresses a fundamental problem called the **Mote-UV problem** that arises in sensing applications [8] and is stated as follows: given the locations of the motes, the amount of data that needs to be transferred from each mote to the base station and an UV, find (i) a communication network among the motes, (ii) a subset of motes referred to as *cluster-heads* that act as relays between the motes and the UV, and (iii) a path for the UV such that

- each mote can use the communication network to transmit its data to one of the cluster-heads,
- each of the cluster-heads is visited by the UV, and,
- the sum of the communication costs involved in transmitting the data from the motes to the cluster-heads and the travel cost of the UV is a minimum.

Refer to figure 1 for an illustration of the mote-UV problem. A cluster-head is responsible for collecting the data from the motes communicating with it and relaying the data to the UV. We allow any two motes to be able to communicate with each other only if they are located within a certain range. The network will mostly consist of stationary motes if the communication costs are relatively cheap in an application; on the other hand, if the communication costs are high and the sensing areas are disjoint, the UV can be used in conjunction with the motes for transferring data.

There are two underlying sub-problems in the mote-UV problem. The first sub-problem deals with synthesizing a communication network that transmits the data from the motes to the cluster-heads and the second sub-problem deals with finding a path for the UV, which in turn transfers the data from the cluster-heads to the base station. These two sub-problems are coupled by the fact that the choice of the motes that will play the role of cluster-heads is not known a priori. In the simplest of the settings where each mote acts as a cluster-head, this problem is a generalization of the classic

1. Mechanical Engineering, Texas A & M University, College Station, TX.

2. Department of Electronics and Communications Engineering, Indraprastha Institute of Information Technology - Delhi, India.

3. Department of Electronic Systems, Aalborg University, Denmark.

4. Faculty of Engineering, University of Porto, Portugal - 4200-365. D. E. Lucani is also with the Instituto de Telecomunicações, Portugal.

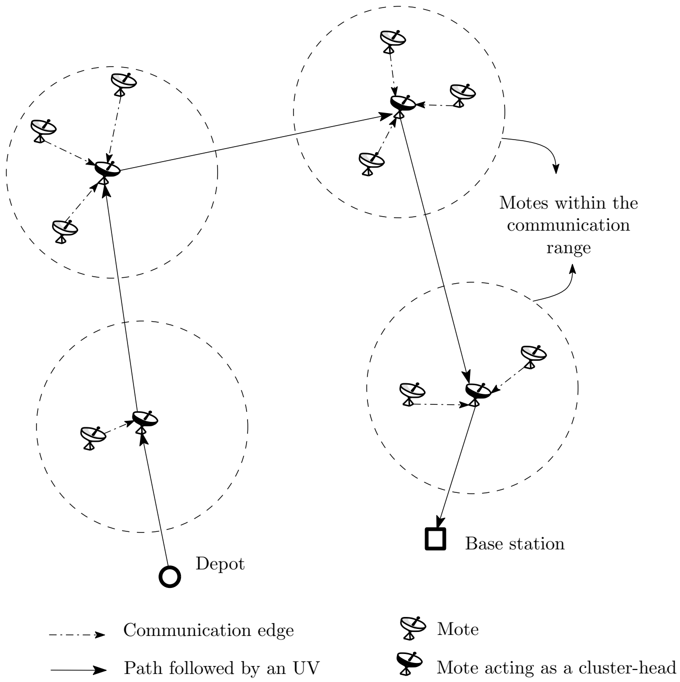

Fig. 1. An illustration of the Mote-UV problem.

Traveling Salesman Problem (TSP) and is NP-Hard [9].

The mote-UV problem is different from the data collection problems addressed in the literature [7], [10], [11]. For example, in [10], the effective communication range of each cluster-head is specified by a disk and the goal is to find a tour such that the UV visits at least one point from the disk of each cluster-head and the length of the tour is a minimum. The authors in [10] pose this problem as a Traveling Salesman Problem with Neighborhoods (TSPN) and present approximation algorithms. Our problem fundamentally differs from the TSPN for the following reason: we select the cluster-heads in order to provide a tradeoff between the communication cost and the routing cost; this feature is not present in the TSPN. In [11], the authors consider a simpler problem where the motes have already been partitioned into clusters. An experimental demonstration of an UV collecting data from a set of motes was discussed in [7].

There are also other data collection problems considered in the literature; however, in [12],[13], the focus is on variants of a one-dimensional data collection problem for UVs while in [14],[15], the focus is on controlling the speed profile of the UVs while assuming there is a pre-defined path specified for each of them.

The following are the contributions of this article:

- We provide a Mixed Integer Linear Program (MILP) to find an optimal solution to the mote-UV problem.
- We develop a rounding algorithm by first solving a Linear Programming (LP) relaxation of the MILP and assigning a subset of motes for the UV to visit. We then use some of the best heuristics available in the literature to solve the resulting TSP and the communication network problem.
- We also develop fast, clustering heuristics to develop feasible solutions to the mote-UV problem.
- All the proposed algorithms are implemented for hundreds of instances to corroborate their performance. The simulation results show that the proposed rounding algorithm produce high quality feasible solutions relatively quickly.

II. Problem Formulation

Let N denote the set of all the nodes corresponding to the motes in the sensor network. Let the nodes s and t denote the depot (initial location) and base station (final location) of the vehicle respectively. Let $\mathcal{N}_c := N \bigcup \{s,d\}$. Let $\mathcal{A}_v$ represent the set of all the directed edges between any two nodes in $\mathcal{N}_c$. The binary variable x_{ij} is used to denote the use of the directed edge $(i,j) \in \mathcal{A}_v$ from node i to node j; that is $x_{ij} = 1$ if the vehicle travels from i to j and $x_{ij} = 0$ otherwise. The choice of whether mote i is acting as a cluster-head or not is determined by the binary decision variable m_i. That is, m_i is equal to 1 if mote i is acting as a cluster-head and m_i is equal to 0 otherwise. In the ensuing discussion, we first formulate the routing and communication constraints, and later present the objective.

A. Route Constraints

A path for the vehicle must start at its depot s, visit each of the cluster-heads exactly once and reach the base station d. We use the multi-commodity flow constraints available for general vehicle routing problems [9] to formulate this path constraint. In this formulation, a path is viewed as a collection of edges through which a distinct unit of commodity can be shipped from the depot to each of the cluster-heads using the vehicle. A commodity can only originate at the depot and must be delivered at its respective cluster-head without accumulating at any of the intermediate nodes. These flow constraints ensure that each cluster-head is visited exactly once by the vehicle and all the cluster-heads are connected to the depot. Suppose f_{ij}^k denotes the commodity corresponding to the k^{th} mote flowing from node i to j. Then, the multi-commodity formulation for specifying a path for the vehicle can be formulated as follows:

$$\sum_{j:(i,j)\in\mathcal{A}_v} x_{ij} - \sum_{j:(j,i)\in\mathcal{A}_v} x_{ji} = \begin{cases} 1 & \text{for } i = s, \\ -1 & \text{for } i = d, \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

$$\sum_{j\in\mathcal{N}_c\setminus\{s\}} (f_{sj}^k - f_{js}^k) = m_k \quad \forall k \in N, \quad (2)$$

$$\sum_{j\in\mathcal{N}_c\setminus\{i\}} (f_{ij}^k - f_{ji}^k) = 0 \quad \forall k \in N, \forall i \in \mathcal{N}_c \setminus \{s\} \text{ and } i \neq k, \quad (3)$$

$$\sum_{j\in\mathcal{N}_c\setminus\{k\}} (f_{kj}^k - f_{jk}^k) = -m_k \quad \forall k \in N, \quad (4)$$

$$0 \leq f_{ij}^k \leq x_{ij} \quad \forall i,j \in \mathcal{N}_c, \ \forall k \in N, \quad (5)$$

$$x_{ij} \in \{0,1\}, \forall (i,j) \in \mathcal{A}_v, \quad (6)$$

$$m_i \in \{0,1\} \quad \forall i \in N. \quad (7)$$

Constraints in (1) state that the path for the vehicle must start at the depot and end at the base station. The multi-commodity constraints in (2)-(4) state that the vehicle must visit any mote that is chosen as a cluster-head.

B. Communication constraints

Let z_{ij} represent the binary variable that specifies if mote $i \in N$ is sending all its data to mote $j \in N$ or not; z_{ij} is 1 if mote i is sending its data to mote j and is equal to 0 otherwise. Then, the communication constraints can be formulated as follows:

$$z_{ij} \leq m_j \quad \forall i, j \in N, \tag{8}$$

$$\sum_{j \in N} z_{ij} = 1 \quad \forall i \in N, \tag{9}$$

$$z_{ij} \in \{0, 1\} \quad \forall i, j \in N. \tag{10}$$

The constraint $z_{ij} \leq m_j$ ensures that mote i can send data to mote j only if mote j is a cluster-head. The constraint $\sum_{j \in N} z_{ij} = 1$ ensures that mote i must send all its data exactly to one of the motes in the collection.

C. Objective

Let us assume that mote $i \in N$ has D_i units of data that must be transmitted to a cluster-head. Let the cost of sending one unit of data from mote $i \in N$ to mote $j \in N$ be denoted by a_{ij}. a_{ij} can be computed in the following way: if two motes lie within the communication range, the communication cost between the motes can be modeled using a first-order radio model [16]. In this model, we assume that the cost of communicating one unit of data between two motes is directly proportional to the square of the Euclidean distance between the two motes. In scenarios where multi-hop communication is allowed, one can construct a communication graph where an edge between any two motes is present in the graph if and only if the motes are within the communication range. As described earlier, a communication cost can be assigned to each of these edges. Any shortest path algorithm can then be implemented on this graph to compute the cost of communicating the data between any two motes. In scenarios where there is no communication path between two motes, the corresponding communication cost will be set to a large positive constant.

The mote-UV problem aims to find a path for the vehicle and a communication network such that the summation of the travel cost for the vehicle and the communication costs is a minimum. This problem is formulated as a Mixed Integer Linear Program (MILP) as follows:

$$\min \sum_{(i,j) \in \mathcal{A}_v} h_{ij} x_{ij} + C_t \sum_{i \in N} D_i \sum_{j \in N} a_{ij} z_{ij}, \tag{11}$$

subject to

Path constraints for the vehicle given by (1)-(7),

Communication constraints given by (10),

$$\sum_{j | (j,i) \in \mathcal{A}_v} x_{ji} = m_i \quad \forall i \in N,$$

where h_{ij} is the cost of traveling from vertex i to vertex j, a_{ij} is the communication cost between vertex i and vertex j, and C_t is the tradeoff constant between the travel cost and communication costs.

III. Rounding Algorithm

This article proposes a general approach for finding feasible solutions for the mote-UV problem based on a Linear Programming (LP) based rounding technique. The LP relaxation of a mixed integer linear program (MILP) is obtained by relaxing all the integrality constraints of the decision variables in the formulation (section II). In the context of the mote-UV problem, the relaxation yields a lower bound to the optimum. The rounding technique involves solving a LP relaxation of the MILP and rounding a set of fractional decision variables to obtain a feasible solution to the MILP. This method has been applied effectively to several NP-Hard problems including the general assignment problem [17], lot-sizing [18], [19], multi-period facility location [20] and network design [21], [22]. The effectiveness of the method relies on the tightness of the bound produced by the LP relaxation.

The rounding algorithm proposed in this article first solves the LP relaxation and rounds some decision variables to either 0 or 1. Specifically, the relaxed variable m_i in the LP relaxation is rounded to either 0 or 1 depending on a given constant ρ. For all $i \in N$, mote i is chosen as a cluster-head (m_i is rounded to 1) if the relaxed variable m_i in the LP relaxation has a value at least equal to ρ. Mote i communicates its data to a neighboring mote or to a cluster-head if the relaxed variable m_i in the LP relaxation has a value less than ρ. This rounding step decouples the communication network synthesis and the path problem for the vehicle. Once the cluster-heads are selected, the communication path from each mote to a cluster-head is solved using a shortest path algorithm. The path for the vehicle is solved using the Lin-Kernighan Heuristic (LKH) [23] which is one of the best heuristics for solving the TSP.

There is one critical step in the rounding algorithm that requires a careful consideration. It is possible that for some values of ρ, the choice of cluster-heads *only* based on rounding does not lead to a feasible solution to the mote-UV problem. This can happen if one cannot find a communication path between each mote and a selected cluster-head. To address this issue, we propose the following procedure in this article: First, we form a communication graph containing the nodes in N where an edge joins two motes in N if and only if the motes lie within the communication range. Next, we find all the connected components of the communication graph. If each connected component has at least one mote i for which $m_i \geq \rho$ in the LP solution, it is easy to check that the proposed procedure will find a feasible solution to the mote-UV problem. If there is any connected component in the communication graph for which there is no mote i such that $m_i \geq \rho$ in the LP solution, at least one of the motes from the connected component is randomly chosen as the cluster-head. For a given ρ, this procedure is guaranteed to provide a feasible solution to the mote-UV problem. The rounding algorithm can be applied for a range of values of ρ and the feasible solution that provides the least cost can be chosen as the output for the rounding algorithm.

IV. Clustering heuristics for the mote-UV problem

In this section, we present fast heuristics that aim to find feasible solutions to the mote-UV problem. The approach here is to decompose the mote-UV problem into subproblems so that the best available search strategies for the sub-problems can be used to solve the mote-UV problem. In particular, we decouple the mote-UV problem into three subproblems as follows:

- *Cluster formation:* Given the motes and their communication range, the first subproblem determines the clusters of motes such that any two motes within a cluster lies within the communication range.
- *Cluster connection:* The second subproblem aims to find the order in which the clusters must be visited and the vertices the vehicle must use to enter and exit in each cluster. We do not restrict the entry and exit vertices of each cluster to be the same.
- *Path Planning inside each cluster:* Given an entry and exit vertex in each cluster, the third subproblem aims to find a path that minimizes the sum of the travel and communication costs within the cluster.

In the following subsections, we discuss the algorithms used for solving each of the subproblems.

A. Cluster formation

In general, for a given set of motes, clusters can be formed using algorithms like k-means clustering, spectral clustering, etc. However, these algorithms do not take the communication range limitations into account. Hence, we developed a tree growing approach to determine a cluster that satisfies the communication range limitations. In this approach, we create a graph of the nodes based on the communication range constraint. Each disconnected subgraph forms a cluster. The number of clusters formed are based on number of disconnected subgraphs in the deployment.

B. Cluster connection

The problem of finding the entry and exit notes in each cluster, and the order of visiting the clusters is a generalization of the One-in-a-set TSP which is NP-Hard. To address this problem, we present two heuristics in this article. The first heuristic is a simple nearest neighbor heuristic while the second heuristic relies on solving a TSP.

1) Nearest Neighbor (NN) heuristic: In the NN heuristic, the mote that is closest to the depot (with respect to the traveling cost) is first chosen as the entry node (s_1) into the cluster containing s_1. Next, s_2 is chosen such that s_2, s_1 belong to different clusters, the cluster containing s_2 does not have a entry node and s_2 is the closest node to s_1. This process is iteratively applied to find the entry node to each of the clusters. Once the entry nodes are chosen, the exit node (say t_i) is chosen such that s_i, t_i belong to the same cluster and t_i is closest to s_{i+1}. For the cluster that is visited immediately prior to visiting the base station, the exit node is chosen such that it is the closest node to the base station.

2) Meta-Heuristic: In this heuristic, the median for each cluster is determined. The nearest node to the median is selected as the source node for that cluster. Using the source nodes of all the clusters, a TSP solution is generated using a swarm optimization meta-heuristic[24]. The route given by the TSP solution will determine the sequence in which the UV must visit the clusters. The exit node of a cluster can be either the source node or a node that is nearest to the source node of the next cluster. Using the source and exit nodes of a cluster, paths are determined using the algorithm giving in the following subsection.

C. Path planning inside each cluster

To determine the travel path within each cluster for the vehicle, we first identify the best k-shortest paths with respect to the travel cost and choose the path that minimizes the total cost for the cluster. Suppose the vertices in a cluster are denoted by V^i. Let s_i and t_i be the entry and exit node of this cluster. The goal is to find a travel path from s_i to t_i such that the sum of the travel cost of the vehicle and the communication cost of all the motes within the cluster is minimized. The k-shortest path algorithm[25] is first applied on the nodes in V^i to find a collection of k shortest paths, $PATHS(V^i)$, with the least traveling cost from s_i to t_i. Given a path from the collection, the cost of communicating the data from a mote $k \in V^i$ not in the path to any cluster-head on the path can be simply computed using a shortest path algorithm. For a given path denoted by P, the proposed heuristic assigns a cost defined as $Cost(P) := Travel(P) + C_t \times Comm(V^i)$ where $Travel(P)$ is the total length of the travel path from s_i to t_i and $Comm(V^i)$ is the sum of the communication costs of all the motes within the cluster. The proposed heuristic then chooses a path P^* such that $Cost(P^*) = min_{P \in PATHS(V^i)} Cost(P)$.

V. Simulation results

This section compares the performance of the proposed algorithms with respect to the deviations of the solutions obtained by the algorithms from the optimum. For the simulations, we considered problems of sizes ranging from 15 motes to 50 motes in increments of 5. For each problem size, 50 random instances were generated and the location of a mote in each instance was randomly chosen in a square area of 1000x1000 units. In addition, each instance had the depot and the base station placed at the origin. All the simulations were run on a Dell Precision $T5500$ workstation (Intel Xeon $E5630$ processor 2.53GHz, 12 GB RAM).

The Euclidean distance between the two locations was chosen to be the travel cost between the locations. The communication cost, or the energy spent in transmitting data between any two motes, was modeled using a first-order radio model[16]. For the simulations, we assumed that each mote has one unit of data to transmit to the vehicle. The

TABLE I

ROUNDING ALGORITHM: AVERAGE DEVIATION FROM OPTIMUM (IN %) FOR DIFFERENT VALUES OF ρ

No. of Nodes	$\rho = 0.4$	$\rho = 0.5$	$\rho = 0.6$
15	2.05	1.14	1.14
20	2.92	1.85	1.85
25	3.75	1.95	1.95
30	5.08	2.76	2.76
35	5.85	2.97	3.02
40	6.24	3.55	3.54
45	7.42	4.49	4.49
50	8.50	5.53	5.59

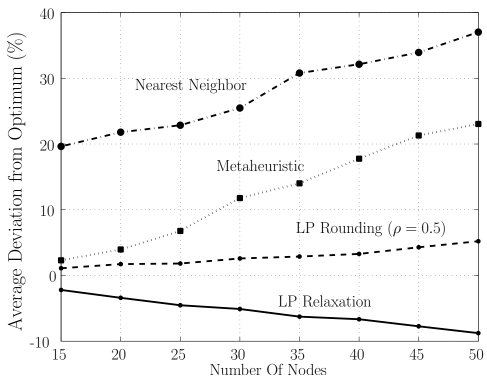

Fig. 2. The average deviation of the solutions (in %) produced by all the proposed algorithms when $C_t = 0.0001$. The LP rounding heuristic produced the best solutions to the mote-UV problem on an average.

communication range of each mote was chosen to be 100 meters. We did simulations for two different values of the tradeoff constant C_t (0.0001 and 0.001). We note here that small values for the tradeoff constant are realistic since the travel cost of a vehicle is typically more costly than the transmission of a small number of packets.

For each of the problem instances, both the MILP and the LP was solved to optimality using IBM ILOG CPLEX 12.0. The LP rounding heuristic was coded in C++ using the Boost Graph Libraries (BGL) and the other heuristics were coded using MATLAB. CPLEX required 7 hours of CPU time for solving most of the instances of the mote-UV

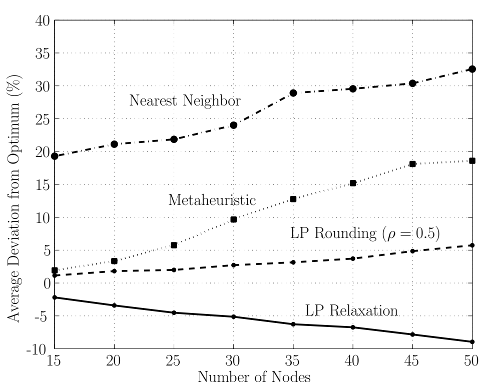

Fig. 3. The average deviation of the solutions (in %) produced by all the proposed algorithms when $C_t = 0.001$. The LP rounding heuristic produced the best solutions to the mote-UV problem on an average.

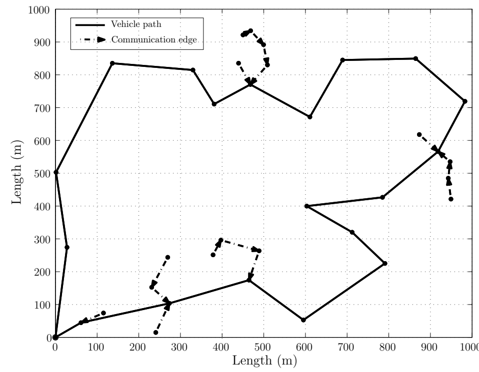

Fig. 4. Optimal solution for an instance with 35 nodes.

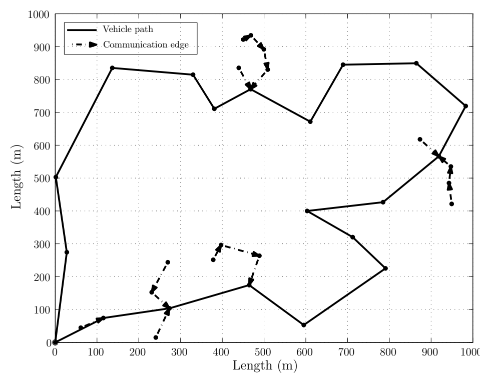

Fig. 5. Solution obtained using the rounding heuristic for the instance with 35 nodes.

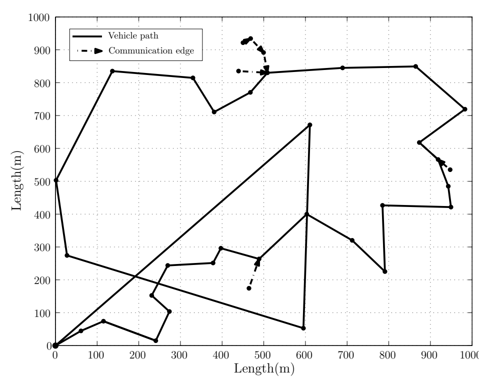

Fig. 6. Solution obtained using the nearest neighbor heuristic for the instance with 35 nodes.

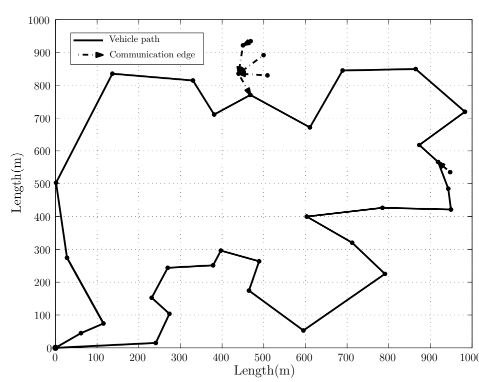

Fig. 7. Solution obtained using the meta-heuristic for the instance with 35 nodes.

problem with 50 motes. On the other hand, the LP rounding heuristic required 2 minutes of CPU time on an average to solve each of the instances with 50 motes.

For a given instance I, the deviation of the solution produced by an algorithm from the optimum is computed as follows:

$$\frac{\text{Cost(Algorithm)}_I - \text{Cost(Optimal)}_I}{\text{Cost(Optimal)}_I} \times 100\%,$$

where $Cost(Algorithm)_I$ is the cost of the solution produced by the algorithm and $Cost(Optimal)_I$ is the optimal cost for instance I.

In the first set of simulations, we compared the performance of the rounding algorithm for different values of ρ. Table I shows the average deviation of the solutions when $\rho = 0.4, 0.5$ and 0.6 and the tradeoff constant $C_t = 0.0001$. A quick glance of the table shows that the value of 0.5 for ρ produced the best results for the rounding algorithm. In the next set of simulations, we compared the performance of the rounding algorithm (with $\rho = 0.5$) with the clustering based heuristics, namely the nearest neighbor heuristic and the meta-heuristic. Figure 2 shows the average deviation (in %) for $C_t = 0.0001$ and figure 3 shows the corresponding results for $C_t = 0.001$. As shown in these figures, on an average, the rounding algorithm performed the best and produced solutions within 5% of the optimum. An optimal solution and the solutions obtained using the heuristics for an instance with 35 motes are shown in figures 4-7.

VI. Conclusion

This article poses a basic data collection problem involving a single vehicle with an objective of minimizing the sum of the travel and communication costs. Three heuristics were presented along with a mixed, integer linear program to solve the problem. While CPLEX required hours of computation time to find an optimal solution, the simulation results show that the LP rounding heuristic found solutions within 5% of the optimum (on an average) in the order of minutes. There are several directions in which the data collection problem and the algorithms can be extended. Future work can consider multiple vehicles with motion constraints and algorithms that provide a priori approximation guarantees.

References

[1] K. Geiser, D. Slack, E. Allred, and K. Stange, "Irrigation scheduling using crop canopyair temperature differences," *Transactions of the American Society of Agricultural and Biological Engineers*, vol. 25, no. 3, p. 689694, 1982.

[2] R. Jackson, S. Idso, R. Reginato, and P. Pinter, "Canopy temperature as a crop water stress indicator," *Water Resources Research*, vol. 17, no. 4, p. 11331138, 1981.

[3] J. A. Curry, J. Maslanik, G. Holland, and J. Pinto, "Applications of aerosondes in the arctic." *Bulletin of the American Meteorological Society*, vol. 85, no. 12, pp. 1855 – 1861, 2004.

[4] C. E. Corrigan, G. C. Roberts, M. V. Ramana, D. Kim, and V. Ramanathan, "Capturing vertical profiles of aerosols and black carbon over the indian ocean using autonomous unmanned aerial vehicles," *Atmospheric Chemistry and Physics*, vol. 8, no. 3, pp. 737–747, 2008. [Online]. Available: http://www.atmos-chem-phys.net/8/737/2008/

[5] T. Zajkowski, S. Dunagan, and J. Eilers, "Small UAS communications mission," in *Eleventh Biennial USDA Forest Service Remote Sensing Applications Conference, Salt Lake City, UT*, 2006.

[6] A. Nawoj, "Energy tradeoffs between collection-based routing protocols in wireless sensor networks," *Technical Report*, vol. TR-03-04, 2004.

[7] O. Tekdas, V. Isler, J. Lim, and A. Terzis, "Using mobile robots to harvest data from sensor fields," *IEEE Wireless Communications*, vol. 16, no. 1, pp. 22 –28, Feb. 2009.

[8] P. B. Sujit, D. Lucani, and J. Sousa, "Bridging cooperative sensing and route planning of autonomous vehicles," *Selected Areas in Communications, IEEE Journal on*, vol. 30, no. 5, pp. 912–922, 2012.

[9] D. L. Applegate, R. E. Bixby, V. Chvatal, and W. J. Cook, *The Traveling Salesman Problem: A Computational Study (Princeton Series in Applied Mathematics)*. Princeton, NJ, USA: Princeton University Press, 2007.

[10] B. Yuan, M. Orlowska, and S. Sadiq, "On the optimal robot routing problem in wireless sensor networks," *IEEE Transactions on Knowledge and Data Engineering*, vol. 19, no. 9, pp. 1252 –1261, Sept. 2007.

[11] F.-J. Wu, C.-F. Huang, and Y.-C. Tseng, "Data gathering by mobile mules in a spatially separated wireless sensor network," in *Proceedings of the 2009 Tenth International Conference on Mobile Data Management: Systems, Services and Middleware*, ser. MDM '09. Washington, DC, USA: IEEE Computer Society, 2009, pp. 293–298. [Online]. Available: http://dx.doi.org/10.1109/MDM.2009.43

[12] R. Sugihara and R. Gupta, "Optimal speed control of mobile node for data collection in sensor networks," *IEEE Transactions on Mobile Computing*, vol. 9, no. 1, pp. 127 –139, Jan. 2010.

[13] ——, "Speed control and scheduling of data mules in sensor networks," *ACM Transactions on Sensor Networks (TOSN)*, vol. 7, no. 1, p. 4, 2010.

[14] S. Smith, M. Schwager, and D. Rus, "Persistent monitoring of changing environments using a robot with limited range sensing," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, 2011, pp. 5448–5455.

[15] S. L. Smith, M. Schwager, and D. Rus, "Persistent robotic tasks: Monitoring and sweeping in changing environments." *IEEE Transactions on Robotics*, vol. 28, no. 2, pp. 410–426, 2012.

[16] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan, "Energy-efficient communication protocol for wireless microsensor networks," in *Proceedings of the 33rd Hawaii International Conference on System Sciences-Volume 8 - Volume 8*, ser. HICSS '00. Washington, DC, USA: IEEE Computer Society, 2000, pp. 8020–. [Online]. Available: http://dl.acm.org/citation.cfm?id=820264.820485

[17] A. P. French and J. M. Wilson, "An LP-based heuristic procedure for the generalized assignment problem with special ordered sets," *Computers & Operations Research*, vol. 34, no. 8, pp. 2359–2369, Aug. 2007.

[18] J. Maes, J. O. McClain, and L. N. Van Wassenhove, "Multilevel capacitated lotsizing complexity and LP-based heuristics," *European Journal of Operational Research*, vol. 53, no. 2, pp. 131–148, July 1991.

[19] J. R. Hardin, G. L. Nemhauser, and M. W. Savelsbergh, "Analysis of bounds for a capacitated single-item lot-sizing problem," *Computers & Operations Research*, vol. 34, no. 6, pp. 1721–1743, June 2007.

[20] T. Melo, S. Nickel, and F. Saldanha-da Gama, "An LP-rounding heuristic to solve a multi-period facility relocation problem," 2009.

[21] P. N. Thanh, O. Péton, and N. Bostel, "A linear relaxation-based heuristic approach for logistics network design," *Computers & Industrial Engineering*, vol. 59, no. 4, pp. 964–975, Nov. 2010.

[22] D. B. Shmoys, C. Swamy, and R. Levi, "Facility location with Service Installation Costs," in *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, ser. SODA '04. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2004, pp. 1088–1097.

[23] K. Helsgaun, "An effective implementation of the Lin–Kernighan traveling salesman heuristic," *European Journal of Operational Research*, vol. 126, no. 1, pp. 106–130, 2000.

[24] M. Hoffmann, M. Mhlenthaler, S. Helwig, and R. Wanka, *Discrete Particle Swarm Optimization for TSP: Theoretical Results and Experimental Evaluations*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2011, vol. 6943, pp. 416–427.

[25] D. Eppstein, "Finding the k shortest paths," *SIAM J. Computing*,

vol. 28, no. 2, pp. 652–673, 1998. [Online]. Available: http://dx.doi.org/10.1137/S0097539795290477