Decentralized K-Means Clustering with MANET Swarms

Ryan McCune, Greg Madey
Computer Science & Engineering, University of Notre Dame
Notre Dame, IN 46617
rmccune@nd.edu

Keywords: Agent-Based Modeling, Swarm Intelligence, Emergent Behavior, Clustering, DDDAS

Abstract

A swarm intelligent system is robust, scalable, adaptable, and can efficiently solve complex problems, all through simple behavior. Inspired by biology, swarm intelligent systems, or swarms, utilize emergence, where simple local behaviors distributed across many agents lead to global phenomena, yielding a whole greater than the sum of parts. But the absence of models that quantify emergence, or the lack of an emergent calculus, has challenged swarm engineering. How simple behaviors and interactions lead to complex phenomena is not well understood, let alone developing such behaviors for problem solving.

A swarm intelligent solution is presented to a computationally challenging problem with quantifiable results in support of future models of emergence. The swarm intelligent Decentralized K-Means Clustering technique is introduced within the context of rechargeable Mobile Ad hoc Networks (MANETs). Through engineered emergent behavior, cluster centroids relocate to minimize the sum of the squared error between sensors and the nearest centroid, similar to K-means clustering. An agent-based simulation is developed to evaluate the technique, demonstrating the sum of squared error is consistently reduced for both supervised and random scenarios.

1. INTRODUCTION

Picture a scenario where many small, stationary, wireless sensors are deployed in a large environment. As limited battery life is a critical performance bottleneck for wireless sensor networks (WSNs), the system employs many wireless charging vehicles (WCVs) that traverse the environment and recharge sensors. WCVs first receive charge from a base station, then search the environment for a sensor, and recharge the sensor once in close proximity. After recharging the sensor, a WCV returns to a base station to recharge itself, before heading out once more in search of sensors. The system includes many base stations, which are mobile and capable or relocating. For the system to efficiently recharge the sensors, the mobile base stations should be placed as to minimize the cumulative distance traveled by WCVs.

However, for some reason, be it wireless interference, security, or privacy concerns, only local communication is possible, meaning information is only be exchanged when a WCV returns to a base station. Non-local communication, such as a remote operator controlling a mobile vehicle, or long-range broadcasts of information, is not possible. A corollary of this constraint is that there is no sense of global positioning, such as coordinates, and the system is unaware of sensor locations. But then how can the WCVs navigate between the sensors and base stations? How can the base stations be intelligently placed without being told where to go?

The presented solution employs swarm intelligence. Swarm intelligent systems are characterized by emergent problem solving capability, where simple behaviors aggregated across many agents give rise to a complex collective behavior. Emergence is often described as a whole greater than the sum of its parts, because the capabilities of the system arise from the synergistic interactions of the agents. For the aforementioned system, simple behaviors are assigned to the WCVs and base stations that result in base stations migrating to efficient locations that minimize WCV travel distance, a technique called Decentralized K-means clustering.

1.1. Decentralized K-Means Clustering

Within the context of rechargeable MANETs, Decentralized K-means clustering utilizes emergent behavior to guide mobile base stations to locations that minimize distance traveled by WCVs. Assuming that sensors are recharged by WCVs traveling from the base station nearest the sensor, sensors can be grouped according to the nearest base station. The process of assigning unlabeled data objects to groups is a well-studied technique in data analysis called clustering [11].

One of the earliest and most frequently utilized clustering algorithms is K-means [16, 10]. K-means partitions n data objects into k disjoint subsets S_j so as to minimize the sum of squared-error

$$J = \sum_{j=1}^{K} \sum_{n \in S_j} |x_n - \mu_j|^2 \tag{1}$$

where x_n represents a data object and μ_j is the geometric center, or the centroid, of the data objects in S_j. Similar to the problem of placing base stations relative to sensor locations, K-means calculates clusters based on the distance between cluster centroids and data objects.

However, K-means is centrally computed by an entity with global coordinate information, in contrast to the constraints of the proposed problem. Often such central algorithms are computationally inefficient, as even K-means is NP-hard [10]. The proposed technique achieves a similar clustering to k-means through emergent behavior, and is called Decentralized K-means. Applications of Decentralized K-means clustering include mobile ad hoc networks.

1.2. Mobile Ad hoc Networks

Micro-electro-mechanical system advancement has enabled development of low-cost, low-power sensor nodes. Mobile ad hoc networks (MANETs) are wireless sensor networks (WSNs) made up of mobile sensor nodes that communicate without a fixed infrastructure or centralized administrator. Nodes can sense and interact with the environment, perform on-board computation, and communicate with other nodes as well as a central base station. Communication may be routed through intermediate nodes. Nodes can be autonomous or centrally controlled. Since nodes are battery powered, system lifetime is a critical consideration for application design.

1.2.1. DDDAS

A related objective is the implementation of Decentralized K-means into the Dynamic Data Driven Application (DDDAS) framework. DDDAS captures the synergistic ability to integrate real-time data into an executing application, while the application also guides the measurement process [3]. Research into the application of DDDAS for the command and control of UAV swarms, a MANET instance, has garnered increasing attention [18]. The non-linear relationship between swarm control parameters and swarm performance makes poor

1.2.2. Wireless Rechargeable Sensor Networks

Breakthroughs in rechargeable lithium batteries and wireless energy transfer have led to Wireless Rechargeable Sensor Networks (WRSNs), where nodes can be wirelessly recharged up to 3 meters [13]. Systems have since been developed that utilize wireless charging vehicles (WCVs), mobile vehicles with high volume batteries that wirelessly recharge sensor nodes [21]. Progress in such systems is evident in the commercial availability of WCVs [28] and the establishment by the Wireless Power Consortium of an international standard for wireless charging interoperability. Well-designed WRSNs offer potentially infinite system lifetimes.

Many wireless charging protocols have addressed different types of WRSNs. Protocols have been designed for a single stationary charger to support many mobile sensor nodes [30], for a single WCV to charge many stationary nodes [21], for multiple WCVs charging many stationary nodes [27], and for sensor nodes to collaboratively charge one other [31]. Often, systems incorporating WCVs borrow concepts from Message Ferrying, a technique for efficient data transfer in sparse WSNs. With Message Ferrying, rather than nodes broadcasting data over costly distances, power is conserved via mobile nodes transporting data between stationary nodes and base stations [33]. One prevailing organization for data ferrying is a three-tiered architecture, with a tier of central base stations, a tier of stationary sensor nodes, and an intermediate tier of mobile transport nodes [22]. Systems have also been proposed that incorporate mobile nodes for both tasks of data gathering and wireless charging [32, 7].

1.2.3. Boomerang Behavior

Whether wirelessly charging or data ferrying, mobile nodes transporting resources between a base station and outer nodes perform a simple, recurrent behavior of seeking out a location of interest, then returning to base. Specifically, a mobile node departs a base station in search of an object or location of interest. Upon detecting the object or location, the node transitions to a return state and heads back to base. The node may stay at the base station for a length of time, or until another signal triggers the node to set out again in search of another object or location. The behavior, deemed Boomerang behavior, is depicted in Figure 1.

The Boomerang behavioral pattern is prevalent in many MANET applications. For example, systems utilizing Unmanned Aerial Vehicles (UAVs) for fire-fighting [24] exercise Boomerang behavior when UAVs search for fires, deploy a fire-retardant payload, then return to base for refill. Moreover, UAV systems delivering medical supplies [5] must replenish from a central repository after providing for the field. Many UAV systems exhibiting Boomerang behavior can be implemented within the DDDAS framework [3]. Almost any system characterized by the three-tiered transport node architecture demonstrates Boomerang behavior.

However, though widespread, Boomerang systems generally require a high degree of coordination, communication connectivity, and computationally-intensive centralized planning, resulting in an complex, energy-intensive system in an

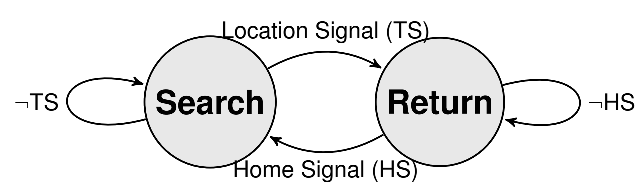

Figure 1: Boomerang behavior of Wireless Charge Vehicles (WCV), where a WCV departs a mobile base station in search of food, and then upon finding food, returns to the nearest mobile base station. The WCV represents the intermediate tier in the three-tiered architecture design pattern (see Figure 2.

often dynamic environment. Application designers, in contrast, strive for systems that are robust, fault-tolerant, flexible, scalable, and computationally undemanding. For such a system, inspiration is drawn from biological swarms.

1.3. Swarm Intelligent Systems

Swarm intelligent systems are characterized by emergent problem solving capability, where simple behaviors aggregated across many agents give rise to complex phenomena, rendering the collective system greater than the sum of its parts [2]. Common examples are observed in nature, such as the coordinated flocking of birds, or ant colonies uncovering efficient paths to a food source. In ant colonies, ants follow pheromone trails deposited in the environment by other ants. Stigmergy describes systems where agent modifications to the environment impact future actions of other agents. Swarm intelligent systems exhibit many advantageous properties, including robustness, flexibility, scalability, and decentralization.

In this paper, we propose a swarm intelligent system for MANETs exercising Boomerang behavior, or Boomerang Swarms. The system implements digital pheromones to facilitate a stigmergic environment for MANETs, resulting in a solution that is fault-tolerant, scalable, and completely decentralized, requiring no (non-local) communication, connectivity, central control or global perspective. Furthermore, unlike any previous biologically-inspired ant system, the Boomerang Swarm implements mobile nests, where nests are guided toward areas of higher food concentration, reducing the cumulative distance between food and nearest nests. An agent-based simulation is employed to evaluate the effectiveness of Boomerang Swarms for Decentralized K-means clustering. An overview of the swarm intelligent system for Decentralized K-means clustering is presented in section 3. The agent-based simulation and evaluation are presented in section 4. Previous work is discussed next in section 2.

2. PRIOR WORK

Much work on the system lifetime performance bottleneck of WSNs and MANETs has focused on methods of energy conservation. An early example is the LEACH protocol [9], where instead of nodes transmitting sensor data directly to the sink, data is routed through local cluster-head nodes of rotating assignment. Other energy conservation approaches include data ferrying, where mobile relay nodes transport data between source and destination nodes in sparse networks [1].

Still, energy conservation methods fail to solve the problem of finite system lifetime presented by battery-powered sensors. Efforts to overcome limited energy storage, other than energy conservation, include energy harvesting, where energy is extracted from the environment through means such as solar panels, wind turbines, or from other sources, including heat, light, radio, or vibrations [20]. Harvesting techniques experienced limited success, in part due to heavy dependency on environmental conditions.

Recent breakthroughs in wireless recharging [13] have turned attention to protocols for WRSNs. Instances of WRSN include a single mobile charger servicing stationary sensor nodes. In such a system, designers often outline a performance metric, prove an NP-hard reduction, and formulate a centralized heuristic or approximation utilizing global information, like location. In [21], a WRSN prototype is developed, where the path of the WCV is planned using a greedy heuristic for the Traveling Salesman Problem. In [23, 29], the WCV traverses a Hamiltonian cycle to optimize the ratio of idle time to the renewable energy cycle. For systems with stationary chargers, [26] investigates optimal node deployment and routing arrangements through reduction of another NP-hard problem. In each case, global information is presumed available, and centralized computation is performed.

While the comparison between mobile charging and data ferrying is often explicit [29, 21, 31], some approaches go so far as to combine the two techniques, forming a system where a mobile vehicle capable of both recharging power and collecting data. [12] explores maximizing the number of sensor-captured events by jointly scheduling a WCV and node duty-cycling. The J-MERDG protocol is proposed in [32], where a WCV, utilizing global knowledge of sensor locations, plans an efficient path between sensor nodes. A joint routing-planning scheme is proposed in [14] that implements energy-balanced and energy-minimum routing. [7] improves upon [32] and [12] by incorporating additional time-varying energy consumption models.

Few WRSNs implement multiple mobile chargers, however. [31] was the first to propose collaborative charging, where mobile nodes coordinate a rendezvous in a one-dimensional system to wirelessly recharge one another. [27] investigates minimizing the total traveling cost of multiple chargers while ensuring no nodes fail, leveraging concepts from Named Data Networking to delivery energy status information. To the best of our knowledge, [17] is the only work to study distributed behaviors of multiple mobile chargers.

In the Boomerang Swarm, the behavior distributed across multiple mobile chargers draws inspiration from swarm intelligent systems. Swarm intelligence is an active subject of interdisciplinary research founded in the study of complex systems. Swarms are capable of self-organization and exhibit decentralized control, rendering the swarm both robust and scalable. The emergent problem-solving ability of ants through pheromone deposits has been developed into an effective meta-heuristic for solving challenging combinatorial optimization problems, such as the Traveling Salesman Problem [6]. An ant-inspired clustering based on ant cemeteries and unrelated to the proposed Boomerang Swarm is presented in [8]. Swarm robotics, and more broadly swarm engineering, are disciplines that apply principles of swarm intelligence to engineering problems. Swarm methods for the command and control of UAVs has received attention.

The central nest foraging problem is an instance of Search-and-Return behavior originally observed in ant colonies [25], and more recently applied within a computer science context [19, 6], where a non-toroidal grid world consists of a nest location and N food source locations. The ant system and agent-based simulation presented in this paper is adapted from [19], with notable differences identified in section 3. In general ant systems, ants begin from the one nest, leave in search of food, and upon finding a food source, become laden with food and return to the nest. For our application, we expand the problem to accommodate M nest locations, where $M \geq 1$, and nests are mobile.

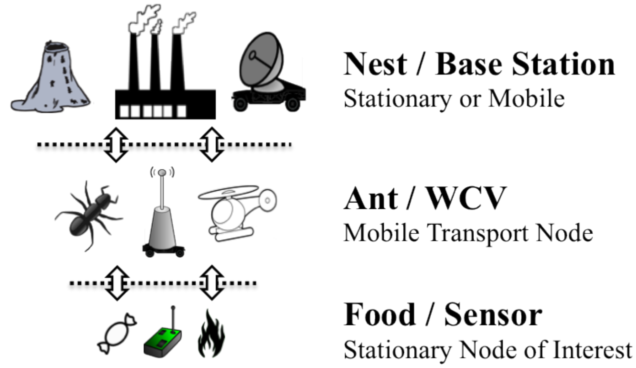

Figure 2: The three-tiered architecture is a recurring MANET design pattern and can be applied to several problems. For WRSNs, the tiers are mobile base stations, WCVs, and stationary sensors. In ant systems, the similar tiers are nests, ants, and food, respectively. The system performs Decentralized K-means clustering, where the mobile base station tier acts as cluster centroids.

3. SYSTEM OVERVIEW

The proposed system is the Boomerang Swarm. The Boomerang Swarm includes three tiers of agents, depicted in Figure 2. Within the context of WRSNs, one tier represents the mobile base stations, an intermediate tier represents the WCVs exhibit Boomerang behavior (see Figure 1), and the stationary sensors are represented by a third tier. Within the context of ant colonies, the tiers represent nests, ants, and food, respectively. Though biological nests are not mobile, the Boomerang Swarm is in general designed for Decentralized K-means clustering, where an intermediate tier exhibiting Boomerang behavior facilitates the relocation of a mobile tier, relative to the other stationary tier. Mobile base stations are initially placed randomly throughout the environment. WCVs begin at base stations and are evenly distributed across the base stations.

3.1. Pheromones

The Boomerang Swarm operates in a stigmergic environment, where indirect communication between WCVs is facilitated by the environment and impacts future actions of the WCVs. At each location, the environment is capable of storing a local pheromone value, represented by a positive integer. For the agent-based simulation presented in section 4, the environment is a grid where each cell possesses a pheromone value. The pheromone value of a particular location can only be sensed by vehicles from adjacent locations, i.e. the 8 Moore neighbors of a cell, made up of the four sides and four diagonal neighbors if the cell. Pheromone values decrease over time at a constant rate, but are increased by a WCV when occupying the location. In application, digital pheromones can be implemented, for example, by small RFID tags dispersed throughout the environment [4], which store an integer value that WCVs can read and write.

In the Boomerang Swarm, WCVs deposit pheromones while in the Search state (see Figure 1). The value of the pheromone to deposit is stored internally by each WCV. The WCVs begin at a mobile base station, where the internal pheromone value is set to a maximum. After each step, the WCV deposits the internally stored pheromone value into the environment, and then the internal pheromone value of the WCV is decreased by a fixed percentage. Thus, a pheromone gradient forms in the environment, where the pheromone concentration is highest around the base stations and decreases further away from base station. WCVs randomly search through the environment, with a lesser probability of moving backwards. Once a sensor is found, a WCV changes to a Return state, ceases depositing pheromones, and follows the path of increasing pheromone concentration, which always leads back to a base station. Once back at a base station, WCVs change state again and the internally stored pheromone value is reset to the maximum. For further details on pheromone implementation, see [19].

3.2. Mobile Base Stations

Previous ant systems and simulations, such as [19], implement the base station tier as a single, stationary agent. The Boomerang Swarm, in contrast, implements many, mobile base stations agents. For base stations to move in a decentralized manner, base stations update location based on the direction of returning WCVs. Mobile base stations detect the returning direction of WCVs, and move one square unit in the direction from which the WCV came. Essentially, mobile base stations take one step closer to the sensor found by the returning WCV. This requires no non-local communication or

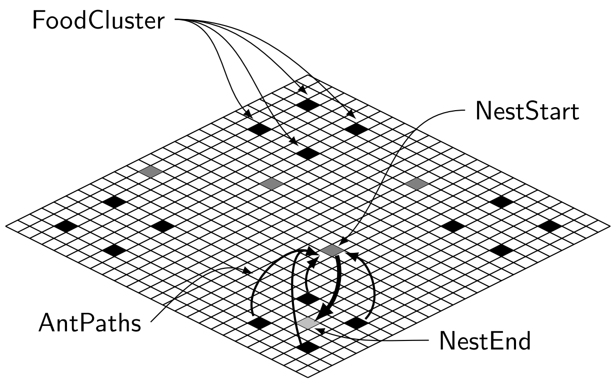

Figure 3: Test scenario with square patterns of food placed in each quadrant. Ideal system behavior would minimize the cumulative distance between food and nearest nests by updating nests to the center of the square.

Parameter	Value
Simulation Duration	10,000 time steps
Environment	120x120, Non-Toroidal
Number of WCVs	100
Number of Mobile Base Stations	4
Number of Sensors	16
Food Placement	Placed, Random
Evaporation Constant	0.999
Update Cut-Down	0.9
Random Movement Probability	0.1
Obstacles	None

Figure 4: Parameters for the Boomerang Swarm agent-based simulation

coordination. To ensure the pheromone gradient always leads back to a mobile base station, after each move the mobile base station updates the new location to a new relative maximum. A general overview of the system, including the relocation of mobile base stations, is depicted in Figure 3. The effectiveness of this simple update mechanism as part of a clustering algorithm is evaluated in section 4.

4. AGENT-BASED SIMULATION

An agent-based simulation of the Boomerang Swarm was developed using MASON, an agent-based modeling toolkit for Java [15]. The simulation adapted previous ant system work provided as part of the MASON toolkit [19]. The implemented environment is a 120 by 120 non-toroidal grid world, with 4 mobile base stations, 16 sensors, and 100 WCVs. Further simulation parameters are outlined in Figure 4.

To evaluate the effectiveness of the Boomerang Swarm for Decentralized K-means clustering, the simulation was used to measure how the cumulative distance of sensors to nearest base stations changes over time. At each time step, the closest base station is determined for each sensor, and the distance is summed over all sensors.

Two different scenarios for placement of the sensors are tested. One scenario, Random Placement, randomly places the 16 sensors throughout the environment. The other scenario, Fixed Placement, places the 16 sensors into a predetermined layout of 4 separate squares, depicted in Figure 3 and Figure 8. Specifically, 4 20-cell by 20-cell squares are placed in each quadrant of the environment, with top-left sensors of each squares located at (20, 20), (20, 80), (80, 20), and (80, 80). Both scenario was run 100 times, each run for 10,000 time steps.

The results are depicted in Figure 5 and Figure 6. Box plots present the distribution of cumulative distances for the 100 simulation runs at every 500 time steps. For both Random Placement and Fixed Placement scenarios, results clearly demonstrate that cumulative distance decreases over time, for all quartiles of the box plot. The random placement of sensors leads to greater distribution of cumulative distances, but also a greater reduction over time.

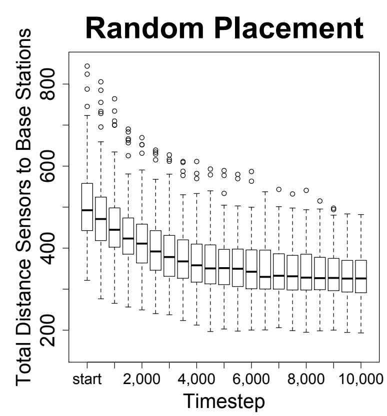

Figure 5: Cumulative distance from sensors to nearest base stations over time for 100 simulations, with sensors randomly placed.

The Fixed Placement scenario can be considered a form of supervised clustering, since groupings are classified. Supervised clustering can be evaluated using the cluster purity measure, which evaluates classification based on the number of correct and incorrect classifications for a given cluster. Purity is scored from 0 to 1, where 1 is considered a perfect clustering. The purity of the Fixed Placement scenario is depicted in Figure 7, showing over half of the 100 Fixed Placement simulations resulted in perfect cluster classification. Moreover, the dotted line in Figure 6 shows the global optimum for mobile base station placement, which was achieved in roughly a quarter of the simulations.

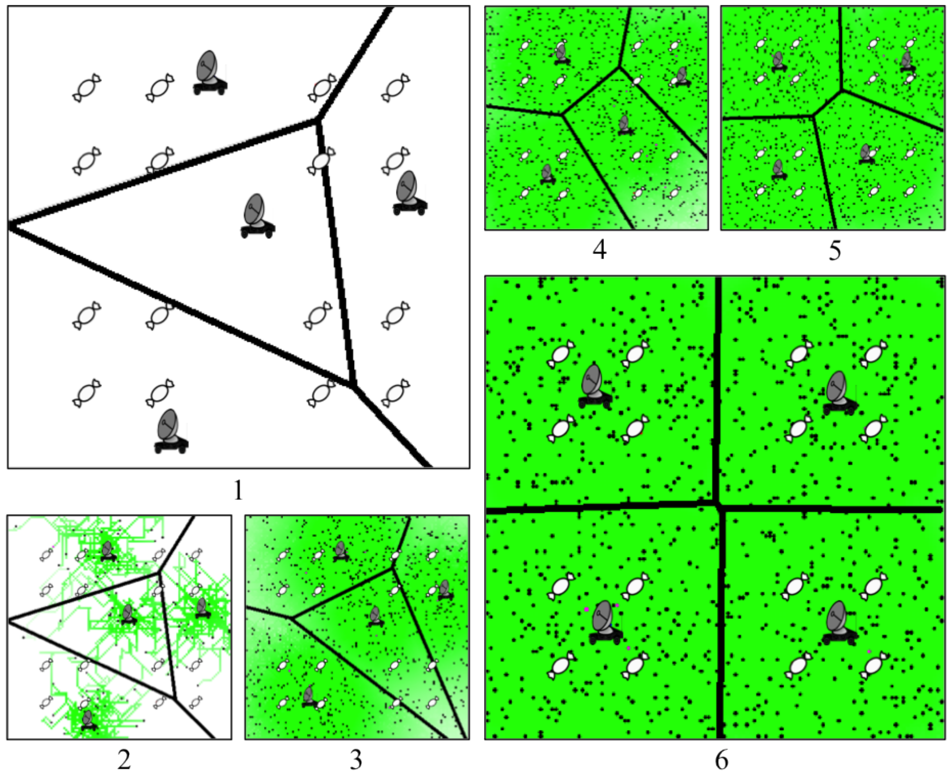

Figure 8: Screenshots from an agent-based simulation of the Boomerang Swarm. Four mobile base stations are initially placed randomly throughout the environment, and 16 sensors, represented by candy, are placed in 4 predetermined squares of 4 sensors each. The hundred black dots are WCVs. At the start, in 1, no pheromones have been deposited. The beginning of pheromone deposits are visible around the base stations in 2, and pheromones are visible in varying shades throughout the environment in 3-6. A Voronoi Tessellation visualization overlays the environment to provide an approximation for which base station a WCV would return to after locating a sensor. Reduction of cumulative distance from food to nests is observed over time.

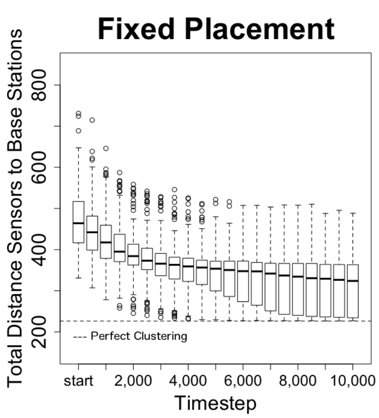

Figure 6: Cumulative distance from sensors to nearest base stations over time for 100 simulations, with sensors placed in 4 predetermined squares of 4 sensors.

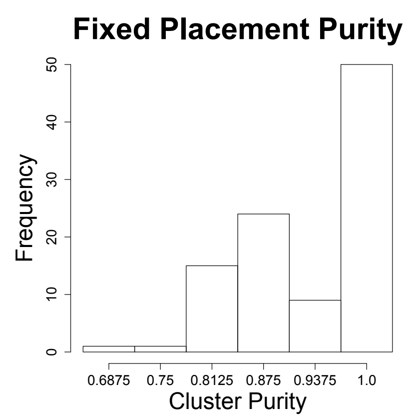

Figure 7: Purity measurement quantifying accuracy of cluster classification, determined at the end of each fixed placement simulation run.

5. CONCLUSION AND FUTURE WORK

A swarm intelligent system called the Boomerang Swarm was presented that collectively performs Decentralized K-means clustering. An agent-based model of the Boomerang Swarm was developed, and the effectiveness of the clustering technique was evaluated for two scenarios, randomly placed sensors and fixed placement into predetermined clusters. Results demonstrate that the technique reduces the cumulative distance between sensors and mobile base stations, though is significantly impacted by the initial random placement of the base stations, much like regular K-means.

The initial presentation of the technique leaves room for and more in-depth evaluation. Further comparisons can be drawn between the Boomerang Swarm and regular K-means by the two methods on the same data sets. Moreover, the impact of initial placement of mobile base stations can be further investigated. Finally, as WCVs follow pheromone trails with a certain probability of randomly moving off course, the effect of the random movement probability parameter on the final clustering can be explored. The random movement probability parameter could be used to control the swarm, and integrated into the DDDAS framework for command and control of UAV swarms.

ACKNOWLEDGMENTS

This work was funded in part by the Air Force Office of Scientific Research under the DDDAS program, as well as the GAANN Fellowship provided by the University of Notre Dame's Computer Science Department.

REFERENCES

[1] ANASTASI, G., CONTI, M., DI FRANCESCO, M., AND PASSARELLA, A. Energy conservation in wireless sensor networks: A survey. *Ad Hoc Networks 7*, 3 (2009), 537–568.

[2] BONABEAU, E., DORIGO, M., AND THERAULAZ, G. *Swarm intelligence: from natural to artificial systems*, vol. 4. Oxford university press New York, 1999.

[3] DAREMA, F. Dynamic data driven applications systems: A new paradigm for application simulations and measurements. In *Computational Science-ICCS 2004*. Springer, 2004, pp. 662–669.

[4] DEMIRBAS, M., ARORA, A., AND MITTAL, V. Floc: A fast local clustering service for wireless sensor networks. In *Workshop on Dependability Issues in Wireless Ad Hoc Networks and Sensor Networks (DIWANS/DSN 2004)* (2004), pp. 1–6.

[5] DOHERTY, P., AND RUDOL, P. A uav search and rescue scenario with human body detection and geolocalization. In *AI 2007: Advances in Artificial Intelligence*. Springer, 2007, pp. 1–13.

[6] DORIGO, M., BONABEAU, E., AND THERAULAZ, G. Ant algorithms and stigmergy. *Future Generation Computer Systems 16*, 8 (2000), 851–871.

[7] GUO, S., WANG, C., AND YANG, Y. Mobile data gathering with wireless energy replenishment in rechargeable sensor networks. In *INFOCOM, 2013 Proceedings IEEE* (2013), IEEE, pp. 1932–1940.

[8] HANDL, J., KNOWLES, J., AND DORIGO, M. Ant-based clustering: a comparative study of its relative performance with respect to k-means, average link and 1d-som. *Design and application of hybrid intelligent systems* (2003), 204–213.

[9] HEINZELMAN, W. B., CHANDRAKASAN, A. P., AND BALAKRISHNAN, H. An application-specific protocol architecture for wireless microsensor networks. *Wireless Communications, IEEE Transactions on 1*, 4 (2002), 660–670.

[10] JAIN, A. K. Data clustering: 50 years beyond k-means. *Pattern Recognition Letters 31*, 8 (2010), 651–666.

[11] JAIN, A. K., MURTY, M. N., AND FLYNN, P. J. Data clustering: a review. *ACM computing surveys (CSUR) 31*, 3 (1999), 264–323.

[12] JIANG, F., HE, S., CHENG, P., AND CHEN, J. On optimal scheduling in wireless rechargeable sensor networks for stochastic event capture. In *Mobile Adhoc and Sensor Systems (MASS), 2011 IEEE 8th International Conference on* (2011), IEEE, pp. 69–74.

[13] KURS, A., KARALIS, A., MOFFATT, R., JOANNOPOULOS, J. D., FISHER, P., AND SOLJAČIĆ, M. Wireless power transfer via strongly coupled magnetic resonances. *science 317*, 5834 (2007), 83–86.

[14] LI, Z., PENG, Y., ZHANG, W., AND QIAO, D. J-roc: A joint routing and charging scheme to prolong sensor network lifetime. In *Network Protocols (ICNP), 2011 19th IEEE International Conference on* (2011), IEEE, pp. 373–382.

[15] LUKE, S., CIOFFI-REVILLA, C., PANAIT, L., SULLIVAN, K., AND BALAN, G. Mason: A multiagent simulation environment. *Simulation 81*, 7 (2005), 517–527.

[16] MACQUEEN, J., ET AL. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability* (1967), vol. 1, California, USA, p. 14.

[17] MADHJA, A., NIKOLETSEAS, S., AND RAPTIS, T. P. Efficient, distributed coordination of multiple mobile chargers in sensor networks.

[18] MCCUNE, R., PURTA, R., DOBSKI, M., JAWORSKI, A., MADEY, G., MADEY, A., WEI, Y., AND BLAKE, M. B. Investigations of dddas for command and control of uav swarms with agent-based modeling. In *Simulation Conference (WSC), 2013 Winter* (2013), IEEE, pp. 1467–1478.

[19] PANAIT, L., AND LUKE, S. A pheromone-based utility model for collaborative foraging. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 1* (2004), IEEE Computer Society, pp. 36–43.

[20] PARADISO, J. A., AND STARNER, T. Energy scavenging for mobile and wireless electronics. *Pervasive Computing, IEEE 4*, 1 (2005), 18–27.

[21] PENG, Y., LI, Z., ZHANG, W., AND QIAO, D. Prolonging sensor network lifetime through wireless charging. In *Real-Time Systems Symposium (RTSS), 2010 IEEE 31st* (2010), IEEE, pp. 129–139.

[22] SHAH, R. C., ROY, S., JAIN, S., AND BRUNETTE, W. Data mules: Modeling and analysis of a three-tier architecture for sparse sensor networks. *Ad Hoc Networks 1*, 2 (2003), 215–233.

[23] SHI, Y., XIE, L., HOU, Y. T., AND SHERALI, H. D. On renewable sensor networks with wireless energy transfer. In *INFOCOM, 2011 Proceedings IEEE* (2011), IEEE, pp. 1350–1358.

[24] SUBRAMANIAM, E., JOSEPH, N., ROSE, A., AND BIL, C. Design of a fire-fighting unmanned air vehicle. In *Aerospace Sciences Meetings*. American Institute of Aeronautics and Astronautics, Jan. 2012, pp. –.

[25] SUDD, J. H., FRANKS, N. R., ET AL. *The behavioural ecology of ants.* Blackie & Son Ltd., 1987.

[26] TONG, B., LI, Z., WANG, G., AND ZHANG, W. How wireless power charging technology affects sensor network deployment and routing. In *Distributed Computing Systems (ICDCS), 2010 IEEE 30th International Conference on* (2010), IEEE, pp. 438–447.

[27] WANG, C., LI, J., YE, F., AND YANG, Y. Multi-vehicle coordination for wireless energy replenishment in sensor networks. In *Parallel & Distributed Processing (IPDPS), 2013 IEEE 27th International Symposium on* (2013), IEEE, pp. 1101–1111.

[28] WATFA, M. K., AL-HASSANIEH, H., AND SALMEN, S. The road to immortal sensor nodes. In *Intelligent Sensors, Sensor Networks and Information Processing, 2008. ISSNIP 2008. International Conference on* (2008), IEEE, pp. 523–528.

[29] XIE, L., SHI, Y., HOU, Y. T., AND SHERALI, H. D. Making sensor networks immortal: An energy-renewal approach with wireless power transfer. *IEEE/ACM Transactions on Networking (TON) 20*, 6 (2012), 1748–1761.

[30] YOON, S. K., KIM, S. J., AND KWON, U. K. Energy relaying in mobile wireless sensor networks. In *Consumer Communications and Networking Conference (CCNC), 2013 IEEE* (2013), IEEE, pp. 673–676.

[31] ZHANG, S., WU, J., AND LU, S. Collaborative mobile charging for sensor networks. In *Mobile Adhoc and Sensor Systems (MASS), 2012 IEEE 9th International Conference on* (2012), IEEE, pp. 84–92.

[32] ZHAO, M., LI, J., AND YANG, Y. Joint mobile energy replenishment and data gathering in wireless rechargeable sensor networks. In *Proceedings of the 23rd International Teletraffic Congress* (2011), ITCP, pp. 238–245.

[33] ZHAO, W., AMMAR, M., AND ZEGURA, E. A message ferrying approach for data delivery in sparse mobile ad hoc networks. In *Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing* (2004), ACM, pp. 187–198.

Biography

RYAN MCCUNE is a PhD candidate in Computer Science at the University of Notre Dame. His research interests include swarm intelligence, agent-based modeling, and multi-agent systems. He is a member of the ACM and IEEE Computer Society.

GREGORY MADEY is a Research Professor in the Department of Computer Science and Engineering at the University of Notre Dame. His research includes topics in bioinformatics, emergency management modeling and simulation, agent-based modeling and simulation, web portals for scientific collaboration, and scientific databases. He is a member of the ACM, AIS, IEEE Computer Society, INFORMS, and the Society for Computer Simulation.