

Exploring Non Uniform Quality of Service for Extending WSN Lifetime

Anna Egorova-Förster
University of Lugano, Switzerland
Email: anna.egorova.foerster@lu.unisi.ch

Amy L. Murphy
University of Lugano, Switzerland & ITC-IRST, Italy
Email: amy.murphy@unisi.ch

Abstract—Interest in extracting and exploiting information from wireless sensor networks has been increasing for several years. Most scenarios focus on collecting and evaluating large amounts of data at a single sink, and many protocols and algorithms have been designed to accomplish this efficiently. However, not all applications require all data, nor do they have a single collection point. Instead, multiple moving base stations collect and exploit information, concentrating on data gathered in their immediate vicinity and requiring less precise knowledge about the data collected far away. While such a scenario can be supported by traditional approaches that present perfect information, they actually require less data, therefore providing a clear opportunity for increasing network lifetime without sacrificing the data quality requirements of the application. This idea of *non uniform quality of service* is the focus of this paper, along with two preliminary protocols demonstrating its potential and motivating future work.

I. INTRODUCTION

Wireless sensor networks have become an interesting and dynamic research field in recent years. As sensor hardware shrinks, new applications emerge ranging from military observation of large-scale battlefields to environmental monitoring and wearable health body sensors.

Research in the wireless sensor domain has concentrated primarily on several application scenarios where data from one or many sources is collected at one destination with high energy efficiency and low latency and overhead. A variety of routing, aggregation, and clustering approaches have emerged from these efforts, some of which incorporate the notion of quality-of-service to ensure accurate data with the required quality is delivered to the destinations. However, the application scenarios they support assume that the destinations, or base stations, need exactly the same data quality from all sources. Nevertheless, some applications can correctly function with non uniform data quality, intuitively allowing energy savings because less data is transmitted.

Consider a disaster recovery scenario in which hundreds or even thousands of sensor nodes are deployed randomly over the affected area. Several fire brigades are sent to the area, each member with his own palmtop to access the sensor network. Based on the reported sensor data, each rescuer independently identifies dangerous regions he can quickly reach.

One option to provide this data is to request all sensors to report data and display a perfectly accurate view of the entire field to each rescuer. While the provided information is sufficient, it is actually more than is required. Instead, each

rescuer only needs to decide where he can go quickly, and if no help is required in his immediate vicinity, he needs to decide a direction to move. Thus, the rescuer needs accurate information from the region immediately surrounding him and only an approximation of the state farther away. While it is straightforward to define such non uniform requirements, challenges arise to support multiple data sinks (rescuers) each requiring a personalized information display, and sink mobility demanding continuous updates to the provided data accuracy.

Exploring the benefits and approaches to provide non uniform data is the main goal of this paper. Our first contribution is the precise definition of non uniform information dissemination in Section II. Related work is outlined in Section III, before we present and evaluate two proof-of-concept protocols in Sections IV and V. These evaluations clearly show the benefits of our approach in terms of potential energy savings, motivating future extensions outlined in Section VI. The paper concludes with a brief summary in Section VII.

II. PROBLEM DEFINITION

As previously introduced, we target a disaster recovery scenario that exploits a WSN of hundreds or thousands of nodes deployed over a large area. The network is assumed to be connected. Each node knows its location, however the global topology is unknown. Throughout this paper we assume the sensors have a fixed location and are homogeneous, all reporting the same types of information, however this is not strictly required. The data collected by the sensors is exploited by a small set of moving rescue workers, or base stations.

While the above matches the typical wireless sensor network scenarios, our work departs in that we assume that the rescue workers require only approximate information about their environment. For example, they may require highly accurate information close to their present location, and only approximate information about distant locations. In other words, the data accuracy required by workers is proportional to the distance between the worker and the data source. Other non uniform quality requirements are also meaningful, e.g., incorporating movement direction to require accurate information in the direction of movement and less accuracy in the movement wake, or adjusting accuracy depending on the density of workers in a particular area. In this paper we focus on the distance-proportional accuracy view as it is both intuitive and demonstrates the challenges of the domain. In any non uniform

data definition, the accuracy is personalized to the location of each worker, thus each should perceive the environment differently and this perspective should be updated as he moves through the region.

III. RELATED WORK

Ideas related to non uniformity have been explored by several research groups. Here we outline a few to place our work in context. The list is not intended to be exhaustive.

This paper represents an extension beyond our own initial exploration of non uniform data dissemination [1] that also presented some preliminary protocols. The primary distinction with respect to the work presented here is the scenario. Previously we assumed all nodes acted as data sinks and the protocols explored various data flooding techniques. Here, our multiple mobile sink scenario offers more opportunities for optimization and distinction from broadcast protocols.

To the best of our knowledge, no comparable efforts exist for achieving non uniformity in WSNs. However, the fish-eye [2] technique from computer graphics has a similar properties, using distance to determine accuracy. This technique inspired Fisheye state routing [3], a MANET routing protocol in which nodes exchange routing tables with frequencies dependent on the distances to the routing table entries. The main difference with respect to our approach is the application of non uniformity itself, namely in the routing tables instead of the data and our assumption of mobile sinks.

Similar non uniformity of data approaches have been introduced in distributed systems [4], [5]. However, neither of these approaches consider energy or CPU processing and both require global knowledge of the static network, thus making them inappropriate for the wireless sensor network domain.

Our work also resembles data clustering and aggregation in WSN. Several approaches [6]–[9] concentrate on creating and maintaining energy-efficient clusters in the network. They have no notion of non uniform cluster sizes or data aggregation techniques and rarely address multiple, mobile sinks. Thus, they cannot be reused for implementing our goals for non uniform quality of service.

IV. ACHIEVING NON UNIFORM DATA DISSEMINATION

Considering our target application scenario, we have identified two promising techniques for achieving non uniform data quality. The first changes the flow of uniformly produced data, e.g. through clustering and aggregation, while the second affects actual data production on the sensors.

A. Non Uniform Data Clustering

The first and most intuitive option to achieve non uniformity of the data in a sensor network is to cluster nodes into non uniform size clusters and aggregate the data inside each cluster before sending it to sink. In other words, data is produced uniformly by all nodes in the network, but the sink does not receive all data. It instead receives aggregated data representing several raw data elements.




Fig. 1. An idealized non uniform clustering. Circles, *bubbles*, represent clusters whose size grows with increasing distance from the sinks (black dots). Clusters are shared among all sinks.

The crucial property for achieving non uniformity of the data is the size of the cluster. Here we assume a simple function that makes clusters bigger farther away from the sink, however as mentioned previously, other functions are possible. To visualize the effect of our simple clustering function, consider each cluster as a *bubble* containing several sensor nodes. Bubbles around a sink are small while those farther away are larger and contain more nodes. With multiple sinks, bubbles should be shared among the sinks as shown in Figure 1. This bubbles analogy takes on additional meaning when considering movement. When a sink moves, bubbles in front of it should break into smaller bubbles, while those behind should merge.

Protocol Definition. Our goal in this paper is to simultaneously provide evidence of the potential benefits of a non uniform data model and demonstrate the feasibility of the approach. Thus we define here a simple, “proof-of-concept” protocol for establishing non-uniform cluster heads for the single, stationary sink scenario.

The first step requires announcement by the sink that wishes to receive data, establishing routing information from each node to the sink. This is accomplished by broadcasting a request from the sink. Upon receipt of this announcement, each node remembers from which node it received the message and its own hop count to the sink, then re-broadcasts the request. Thus, each node eventually receives the request, knows how many hops it is away from the sink, and which node(s) it can use as uplink-nodes for routing the data to the sink. The result is a routing tree rooted at the sink.

After this initialization, each source (typically all nodes) starts gathering data and sending it through one of its uplink nodes. However, before sending the data, it checks whether it is a cluster head or not. For our initial evaluation, we base this decision only on the number of hops to the sink:

$$clusterhead = \begin{cases} true & , hops \text{ is a power of } 2 \\ false & , otherwise \end{cases}$$

If a node *is not* a cluster head, it sends all data it receives (either its own or from other nodes) through the first uplink node. If *it is* a cluster head, it does not immediately forward the




Fig. 2. A sample network with non uniform data clusters. The sink and the cluster heads are circled, two clusters are highlighted with dashed lines.

data. Instead, it waits for some predefined, configurable period, gathers all the data received in this time (including its own), aggregates them according to an aggregation function (to be discussed shortly), marks the new data packet as aggregated, then forwards it to the sink. Already aggregated data packets are not aggregated again, but forwarded directly.

Figure 2 shows a sample network for aggregating data as just described. Uplink connections form a tree rooted at the sink and circled nodes are identified as cluster heads that aggregate data between themselves and the next cluster heads deeper in the tree. Intuitively, the size of the clusters grows with increasing distance from the sink, due to the power-of-2 function used for establishing cluster-head identity. For example, in Figure 2 node 12 aggregates data from nodes 12, 21 and 17 while node 36 aggregates data from all nodes in its subtree: 36, 27, 10, 50 etc.

B. Non Uniform Data Production

A second approach for implementing non uniformity of data in a sensor network is to change the production of the data rather than change its flow in the network. For example, a simple approach is to have only half of the nodes produce data. Alternatively, nodes farther from the sink could produce data less frequently than close nodes. A simple protocol achieving the later idea is described below.

Protocol Definition. As before we assume each sink broadcasts its request for data, but instead of simply informing nodes about their hop count from the sink, the request also defines the *data request function* to be applied at each node. This is a function over the parameters of the request, for example including the data production frequency as in:

$$f(hops) = \begin{cases} \text{FREQ} * hops & , hops \leq 10 \\ 0 & , hops > 10 \end{cases}$$

The effect of this function is to lower the frequency of data production farther away from the sink. If the original frequency is once per second, one-hop nodes will produce data every second, two-hop nodes will produce data every two seconds, etc. Nodes more than ten hops away will not produce data. The overall result is a reduction in the traffic between all sources and the sink.

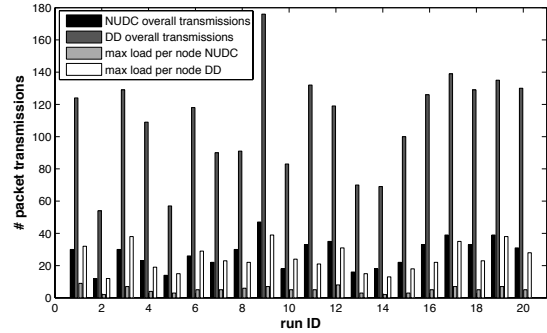


Fig. 3. Energy comparison between DIRECTEDDIFFUSION and NUDATA-CLUSTERING over 20 different topologies. The first two bars indicate the number of transmissions in the whole network, the last two show the load of the maximum loaded nodes.

V. PRELIMINARY EVALUATION

The previously presented proof-of-concept protocols, termed NUDATACLUSTERING and NUDATAPRODUCTION, intuitively reduce data rate. Here we provide a numerical evaluation in comparison to a simple DIRECTEDDIFFUSION technique, similar to the “one phase pull” protocol proposed in [10], in which gradients are built at all nodes towards each of the sinks and all network nodes submit their (non-aggregated) data to the sinks.

All simulations are performed in MATLAB with 50, randomly deployed nodes. Influences from the physical and MAC layers are ignored for this early evaluation, and radio range is simulated with a constant communication radius. The scenarios have one sink (randomly chosen) and all topologies are guaranteed to be connected. Figure 2 represents one scenario we experimented with using 50 nodes in a 100x100 area with a radio range of 25. Our simulation code is available at our website (<http://www.inf.unisi.ch/projects/mics>).

A. Non Uniform Data Clustering

Our first NUDATACLUSTERING experiments use the power-of-two clustering presented earlier and our primary evaluation interest is in the energy consumed for data forwarding. We consider two measures, first the total amount of data transmitted during a single round of data production, and second the maximum amount of data transmitted through a single node. The former reports the energy consumption across the whole network while the later is related to the time to the first node failure. Figure 3 shows results with 20 different initial networks. It clearly shows that aggregating data dramatically reduces overall network cost, a fact that implies longer network lifetime.

These reductions in energy expenditure come at the cost of a loss in data accuracy at the sink. The ability of the sink to reconstruct the data in a meaningful manner depends mostly on the data aggregation function. Here we both average the data collected at the cluster head and report the coordinates of this value as an average of the coordinates of the original data. Figure 4 shows how the reconstruction of the sensor field data




Fig. 4. Reconstruction of sensor data at the sink. The left figure shows the data *after* aggregation of NUDATACLUSTERING and the right without aggregation from DIRECTEDDIFFUSION.




Fig. 5. Energy comparison between DIRECTEDDIFFUSION and NUDATA-APRODUCTION over 20 different topologies, for 100 sec. The first two bars give the number of transmissions in the whole network, the last two give the loads of the maximum loaded nodes.

is done at the sink. The right figure shows the reconstruction with perfect data without aggregation while the left uses only virtual aggregated items. The X and Y coordinates identify the sensor node position and the Z axis represents the data itself, e.g. temperature readings. The reconstructed data fields are extremely similar, implying both that the aggregation function is suitable for this kind of data and its accuracy is not severely impacted by aggregation.

Another important property to consider is the overhead our protocol imposes. Because NUDATACLUSTERING assigns cluster heads during the flooding of data interests, it does not increase the overhead with respect to DIRECTEDDIFFUSION. On the other hand, the latency of the data increases because the cluster head waits to receive multiple values before forwarding the aggregated data. Latency can be tuned by setting the duration of the wait time, yielding corresponding changes in the overall data transmission cost.

B. Non Uniform Data Production

Our second approach, NUDATAPRODUCTION, changes data production at each node according to the earlier function that reduces data production frequency farther from the sink.

As before, our primary goal is to manage the energy production, therefore Figure 5 again presents the total and




Fig. 6. Percentage energy improvement over DIRECTEDDIFFUSION by our protocols for various network sizes. Each point represents an average over 30 different topologies.

maximum single node transmissions. The simulation runs for 100 seconds, using the data production formula from Section IV-B. The highest data production frequency is once per second, meaning nodes one hop away from the sink produce data every second.

There is no difference in the latency or the accuracy of the delivered data, as no aggregation takes place. However, data freshness is affected because distant nodes send data infrequently.

C. Comparing both techniques

As a final consideration, we show the behavior of both protocols as the size of the network increases. Figure 6 shows that, in comparison to DIRECTEDDIFFUSION, NUDATACLUSTERING maintains approximately the same energy savings as the network grows. This is expected because although the clusters far from the sinks increase in size, the communication inside the clusters also increases, thus keeping the relative energy requirements stable. In contrast, NUDATAPRODUCTION causes nodes far away from the sink to produce less data, dramatically reducing the overall amount of data transmitted. Thus, in comparison to DIRECTEDDIFFUSION, the attainable benefits from reducing data that travels multiple hops to the sink grows as the network grows.

VI. PROTOCOL EXTENSIONS

The primary result of our simulations is the positive proof of concept that our ideas have potential applicability in an energy-restricted wireless sensor network. Therefore, the next step is to extend these preliminary protocols in several directions.

A. Features

Thus far the scenario for both the data clustering and data production approaches includes only a single, stationary sink. However our eventual target scenario requires support for multiple, mobile sinks.

Mobile Sinks. We are currently exploring two basic approaches to address mobile sinks, one global and one local.

In the global scenario, the distributed data generated as a result of a request is considered soft state that expires after

a given time period. This requires each sink to periodically renew the request, but also trivially allows the request to change each period. Each time the sink renews the request, it floods the network again, resetting the whole system.

Instead, a locally-restricted approach constantly maintains routing and clustering state. As the sink moves, it triggers routing and clustering updates only where applicable, e.g. in a restricted area around the sink. In reference to the *bubbles* idea presented earlier, this corresponds to breaking big bubbles and merging small bubbles when the sink correspondingly moves in and out of a region.

Multiple Sinks. While both of our current approaches can trivially be extended to maintain distinct state for multiple data sink requests, further increases in network lifetime can be achieved by sharing the aggregation points or merging the data sending functions for multiple sinks. Clearly coordination is required to merge the state from different requests, thus demonstrating the trade-off between the possible savings from sharing and the cost to determine how to share.

B. Approaches

As already noted, the protocols of the previous section are simply a proof of concept to demonstrate the potential benefits of non uniform information dissemination. The protocols themselves are simple, straightforward approaches that can easily be improved through a variety of techniques.

Combining Clustering and Production Functions. So far we have described two approaches, one that defines aggregation points and the other that manages the data production function. We also intend to consider a combination of the two that essentially elects a single node inside each cluster to report data, and turns off data production at all other nodes. Rotating the reporting node can likely further improve results. Overall, while we expect that efficiency will increase, this will come at the cost of determining and rotating the single data producer.

Learning. One approach for solving the general non uniform data dissemination problem is through learning of the best nodes to serve as cluster heads or the best routes for data to take back to the source. Our prior research [11] explores ideas of learning the best routes from a single data source to multiple sinks, using local feedback information to find the optimal results. Our idea is to apply similar feedback techniques locally within the clusters, and globally to adapt as the sinks move and the flow of data changes.

The cost of such an approach is the exploration phase which may make non-optimal selections while searching for the best available solution. By minimizing the duration of the exploration phase, the additional overhead for learning is amortized by the lower cost results.

Exploiting Node Location. Throughout our discussion we have assumed that nodes know their own location, but make all decisions without taking this into consideration. We intend to exploit location information in the learning phase, for example refining the choice of clusters to eliminate extensive overlap in the coverage areas of adjacent clusters. Precisely how to use location depends on the approach, but both clustering

and adjusting the data production can be improved using this additional information.

VII. CONCLUSION AND FUTURE WORK

The main goal of this paper is to define non uniform quality of service for WSN and to explore the opportunities it raises for prolonging network lifetime. As such, we have identified the main challenges and advantages and presented two proof-of-concept protocols. Our initial evaluations of these very simple approaches show that, in general, the concept of non uniform data quality has enormous potential to extend network lifetime over conventional approaches. This clearly motivates future work to extend and enhance our initial protocols to capture all features required by applications that can support non uniform data quality.

ACKNOWLEDGMENT

The work described in this paper is supported by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322.

REFERENCES

- [1] S. Tilak, A. Murphy, and W. Heinzelman, "Non-uniform information dissemination for sensor networks," *Proceedings of the 11th International Conference on Network Protocols (ICNP03)*, Atlanta, GA, USA, November 2003.
- [2] L. Kleinrock and K. Stevens, "Fisheye: A lenslike computer display transformation," University of California, Los Angeles, Tech. Rep., 2004.
- [3] G. Pei, M. Gerla, and T.-W. Chen, "Fisheye state routing in mobile ad hoc networks," in *Proceedings of the ICDCS Workshop on Wireless Networks and Mobile Computing*, 2000, pp. D71–D78.
- [4] P. Yalagandula and M. Dahlin, "SDIMS: A scalable distributed information management system," in *Proceedings of the ACM SIGCOMM Conference, Portland, Oregon*, 2004.
- [5] P. Gross and G. Kaiser, "Automated component interconnection for scaling scale-resistant services," in *Proceedings of the 21st IEEE/ACM International Conference on Automated Software Engineering, Tokyo, Japan*, September 2006.
- [6] H. Luo, F. Ye, J. Cheng, S. Lu, and L. Zhang, "TTDD: Two-tier data dissemination in large-scale wireless sensor networks," *Wireless Networks*, vol. 11, no. 1-2, pp. 161–175, January 2005.
- [7] W. Rabiner-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, Washington DC, USA*, 2000.
- [8] O. Younis and S. Fahmy, "Distributed clustering in ad-hoc sensor networks: A hybrid, energy-efficient approach," in *Proceedings of the IEEE Conference on Computer Communications, Hong Kong*, 2004.
- [9] J. N. Al-Karaki, R. Ul-Mustafa, and A. E. Kamal, "Data aggregation in wireless sensor networks - exact and approximate algorithms," in *Proceedings of the Workshop on High Performance Switching and Routing, Phoenix, AZ*, 2004, pp. 241–245.
- [10] F. Silva, J. Heidemann, R. Govindan, and D. Estrin, "Directed diffusion," in *Frontiers in Distributed Sensor Networks*, 2003, pp. 573–596.
- [11] A. Egorova-Förster and A. Murphy, "A feedback enhanced learning approach for routing in WSN," Faculty of Informatics, University of Lugano, Tech. Rep., 2006.