

# Poster Abstract: Spatial Average of a Continuous Physical Process in Sensor Networks

Saurabh Ganeriwal

Chih Chieh Han

Mani B. Srivastava

Networked and Embedded Systems Lab (NESL), University of California Los Angeles

56-125B Eng. IV, UCLA EE Dept., Los Angeles, CA 90095

{saurabh, simonhan, mbs}@ee.ucla.edu

## ABSTRACT

Wireless ad-hoc sensor networks have caught the fancy of many researchers through the world in a very small span of time. Although notable progress has been made in several key areas, several researchers falter in the way they look at sensor networks i.e. just another kind of wireless ad hoc networks, albeit one composed of energy constrained nodes. A key aspect that makes sensor networks different from traditional networks is their strong link to the physical world. We develop this perspective by studying an interesting class of sensor network applications called *aggregation applications*.

We argue that when a user queries for aggregates like maximum, average etc., he is interested in knowing the statistics of the underlying physical process. In general, this is not equivalent to the statistics gathered from a bunch of nodes. We introduce a distinction between aggregates calculated over a region and aggregates calculated over a discrete set of nodes. We classify them as spatial and nodal aggregates respectively. Till now, researchers have followed the approach of doing nodal aggregation in response to every user query in sensor networks. For a continuous physical process and a random deployment of sensor nodes, which are a norm for sensor networks, the conventional approach of doing nodal aggregation produces inaccurate results.

We verify our claim by studying a specific aggregation function namely *the average over a region (spatial average)* in detail. We use a voronoi-based approach to propose an algorithm for calculating spatial average in sensor networks. The algorithm can work in two modes – centralized and distributed. The energy and accuracy tradeoff associated with each of the two modes is analyzed in detail. The efficacy of the proposed algorithm is tested over a wide variety of physical processes including real precipitation data of South America acquired over a span of past 50 years. We will show that our algorithm gives a 2-8 times performance gain as compared to the conventional approach of doing nodal aggregation. We further explore the robustness of our algorithm to link failures and channel impairments.

Although not 100% accurate, our approach is a simple, efficient and a practical solution for calculating the spatial average. We have implemented our algorithm on Berkeley motes. Our first prototype implementation occupies less than 12K flash ROM and consumes less than 1.6K run-time memory. Our algorithm can be easily integrated with available frameworks for doing aggregation in sensor networks. For a minimal cost, our

algorithm can be used to provide a spatial aggregation service on a network of motes.

## Categories and Subject Descriptors

C.2.4 [Computer Systems Organization]: Computer Communication Networks – *Distributed Systems*.

## General Terms

Algorithms, Design, Performance, Verification.

## Keywords

Sensor Networks, Aggregation, Average, Voronoi.

## 1. INTRODUCTION

Similar to other distributed computing systems, aggregation applications [1] forms one of the most important areas of research in sensor networks. However unlike traditional distributed computing systems, in sensor networks there exists a strong coupling to the physical world. The researchers till now have not appreciated this link between the physical world and the sensor network. It is important to realize that in sensor networks, the user is not interested in values from a bunch of sensor nodes but is instead trying to learn some parameters of the underlying physical process. This physical process can be discrete for example sensor network might be monitoring its own health such as battery lifetime, connectivity etc. However, most of the time the process is going to be continuous such as monitoring temperature, precipitation, levels of CO<sub>2</sub> in air etc. In such scenarios, aggregates will be defined over regions rather than a set of nodes. Thus a typical aggregation query for sensor networks will be of the type “Give me the average temperature in this room” rather than “Give me the average temperature over the nodes  $N_1, N_2, \dots, N_k$ ”. We classify these two types of aggregates as spatial and nodal aggregates respectively.

In case of a uniform deployment of nodes, it can be easily observed that the two aggregates will yield similar results. However, sensor networks are envisioned to operate largely unattended, often in environments where the access cost of deploying and maintaining the nodes is high. In these inhospitable regions, the network designer would have little control over the deployment of the network. Thus an irregular deployment of sensor nodes is the assumed norm for these networks. Even if the nodes are deployed uniformly at the onset of the network, as time progresses, nodes are going to die randomly, resulting in a non-uniform network. For such arbitrary deployment of nodes, there is a need of treating spatial aggregation differently from nodal aggregation.

Realizing the inadequacy of nodal aggregates, there are broadly two approaches to address the problem of spatial

aggregation in sensor networks. The first approach is to reconstruct the whole physical process from these discrete values. However this process is computationally intensive and is not feasible on energy constrained sensor nodes. The second approach is to somehow abstract the missing link of “space” between the continuous process and the discrete set of node values. We follow this approach and provide a voronoi-based solution for calculating spatial aggregates in sensor networks. To the best of our knowledge, nobody has yet tried to appreciate this link between the discrete sensor network and the continuous physical process being monitored by it.

## 2. SPATIAL AGGREGATION

We classify the type of aggregates which are calculated over a continuous region  $R$  and a set of nodes,  $S(t)$ , as *spatial and nodal aggregates* respectively. When faced with a query of calculating the aggregate over a continuous region  $R$ , the conventional approach is to form a subset  $S(t)$ ,  $S_R(t)$ , containing the values from nodes lying in region  $R$ . The final answer is obtained by calculating the nodal aggregate over  $S_R(t)$ . This approach misses the notion of “space”. We try to capture this missing link by appropriately weighing the contribution of a sensor node by the area covered by it. We calculate this area by doing a voronoi tessellation [2] of the sensor terrain.

We have proposed two different modes of our algorithm – centralized and distributed, which corresponds to the way in which the voronoi cell is calculated for every sensor node. In the centralized mode, the algorithm can be run in two different ways – *centralized-p* (periodic) and *centralized-s* (snapshot), depending on the frequency of the aggregation query. We have analyzed the accuracy (error in the calculation of the voronoi cell area) and complexity (energy spent) associated with all the versions of our algorithm.

## 3. SIMULATION RESULTS

In this paper, we justify the distinction between nodal and spatial aggregation by studying a specific aggregation function, the *spatial average*. We have implemented the algorithm in NESLsim, a PARSEC based simulation platform for sensor networks. The efficacy of the proposed algorithm was tested over a wide variety of fabricated physical processes. Our algorithm outperforms the conventional approach by a factor of 2-8. We have verified the robustness of our algorithm to channel impairments and link failures. Finally, we tested the

performance of our algorithm on real precipitation data of South America acquired over a span of past 50 years. Figure 1 plots the average and maximum relative error. All the versions of our algorithm give a 2-5 times better performance than the conventional approach.

## 4. IMPLEMENTATION ON MICA NOTES

We have implemented the distributed version of our algorithm, where every sensor node calculates its own voronoi cell, on Berkeley Mica Motes. Our prototype implementation occupies 11496 bytes flash memory and consumes 1570 bytes run-time memory when computing the voronoi cell for at most 20 neighbors. We have also evaluated energy-efficient operating regions for each version of our algorithm on Berkeley motes. An a-priori knowledge about the number of sensor nodes and the period of the aggregation query can help the network designer to choose the optimal version (from the perspective of energy consumption) of our algorithm. We have encapsulated our algorithm with a very simple API that computes the area enclosed by voronoi edges generated with sensor node’s neighbors. Our algorithm can be easily integrated with existing frameworks for in-network processing such as TAG to provide a spatial aggregation service on a network of motes.

We are currently working towards extending the framework of spatial aggregation to other aggregate functions such as histogram, isotherms, median etc.

## 5. ACKNOWLEDGEMENTS

The abstract is based upon work supported by the National Science Foundation (NSF) and by the Office of Naval Research (ONR) under the AINS Program. The views expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF or the ONR.

## 6. REFERENCES

- [1] A. Boulis, S. Ganeriwal, M. B. Srivastava, “Aggregation in sensor networks: an energy-accuracy tradeoff,” *IEEE workshop on Sensor Network Protocols and Applications (SNPA)*, 2003.
- [2] F. Aurenhammer, “Voronoi Diagrams – A Survey Of A Fundamental Geometric Data Structure,” *ACM Computing Surveys* 23, pp. 345-405, 1991.




Figure 1: Performance of spatial aggregation approach on real precipitation data