

SPEED: Scalable Protocols for Efficient Event Delivery in Sensor Networks

Tassos Dimitriou¹, Ioannis Krontiris¹, Fotios Nikakis¹, and Paul Spirakis²

¹ Athens Information Technology, Markopoulo Ave.,
190 02 Peania, Athens, Greece
{tassos, ikro, fnik}@ait.gr

² Computer Technology Institute, Greece
spirakis@cti.gr

Abstract. One of the most eminent problems in sensor networks is the routing of data to a central destination in a robust and efficient manner. In this work we propose a new scalable protocol for propagating information about a sensed event towards a receiving center. Using only local information and total absence of coordination between sensors our protocol achieves to propagate the sensed data to a receiving center by activating only those nodes that lie very close to the optimal path between the source of the event and the destination, resulting in low activation of the network's sensors. Thus the protocol is very energy efficient. Furthermore, our protocol is robust as it manages to propagate the information even when sensors fail with certain probability.

1 Introduction

Sensor networks[1] have attracted much scientific interest during the past few years. These networks use hundreds to thousands of inexpensive wireless sensor nodes over an area for the purpose of monitoring certain phenomena and capture geographically distinct measurements over a long period of time.

These networks differ from wireless ad hoc networks in the sense that their nodes are characterized by limited resources such as storage, computational and communication capabilities. The power of sensor networks, however, lies exactly in the fact that their nodes are so small and cheap to build that a large number of them can be used to cover an extended geographical area, gather information in-site and process it in parallel enabling an accurate and reliable monitoring process that was previously hard or infeasible to get. Due to the limited resources available to sensors, expensive routing protocols, costly flooding mechanisms, or complex algorithms that don't scale to large number of nodes cannot be used. Furthermore, random distribution of nodes in the physical environment, node failure during their deployment and dynamic change of nodes' power supply make the design of communication protocols a very challenging task.

In this work we focus on the efficient propagation of a sensed event towards some receiving center, assuming an event-driven data delivery model[6]. The need for communication between a regular sensor (the source) and some base

station (called the destination or the sink) can arise at any time, possibly triggered by unexpected changes in the environment. It is exactly this change in the environment (i.e. a fire, a person entering a restricted area, etc.) that we feel it is important to reach the base station as quickly as possible without depleting the network from its resources through the use of complicated protocols.

There is an abundance of routing protocols for sensor networks (see for example [3] and [4] for two distinct representatives) that attempt to optimize a variety of different measures including efficiency, robustness, number of activated particles, etc. In our setting, however, where sensed data need to be sent to a receiving center, these approaches seem to be somewhat of overkill as we only need to send a single message or packet back to the base station. Taking into account the small communication throughput and the limited memory and computational capabilities of sensor networks a simple flooding approach seems to be the best alternative.

Flooding is the most computationally efficient protocol due to its simplicity as every node broadcasts every new incoming packet. Therefore, data are bound to reach their destination, assuring correctness, and the protocol is immune to node failures, assuring robustness. Although this protocol can be integrated even in the most simplistic implementations of sensor nodes, it is extremely energy consuming as all nodes must receive and transmit the message at least once. Gossiping or wandering approaches [2] seem to alleviate this problem, at the cost, however, of increasing path lengths or failing to reach destination.

Our Contribution: In this work we propose a new light-weight protocol that tries maximizing efficiency and minimizing energy consumption by favoring certain paths of local data transmission towards the sink. Just like flooding, the protocol is very easy to implement, even by the simplest of sensor nodes, as it only requires nodes to decide whether or not to forward the message to their neighbors. Unlike flooding protocols, however, it avoids depleting the network from its resources by restricting the nodes that receive and hence retransmit the message. The mechanism that controls this propagation of information is very simple; every node decides to retransmit according to the value of a single parameter that essentially captures the node's distance from the optimal path between source and sink. By carefully tuning the retransmission threshold one obtains a trade-off between robustness (the fraction of times the message reaches the destination) and load incurred in the network (measured in terms of the ratio of activated nodes).

2 Description of the Protocol

In what follows we assume that the number of sensor nodes in the network is N and there is a single destination point D that represents the center where data should be sent. We denote the node that sensed the event by S . Our focus is the design of a simple algorithm that solves this communication problem between S and D in an efficient and robust way by adhering to the following design principles that are particular to sensor networks: (i) low energy requirements

and utilization of the sensors, (ii) scalability and use of local information only, and (iii) robustness under node failures.

The goal of the algorithm is to activate only those sensors that lie very close to the optimal line (in terms of hops) between S and D . The thinner this zone of transmitting nodes is, the more efficient will be the protocol in terms of the ratio of activated nodes. The algorithm is shown below:

Angle Protocol

Let P be the sensor that needs to decide whether to forward a received message or not. Denote by P_{prev} the sensor that transmitted the message to P and by $\omega = (\angle P_{prev}PD)$ the angle defined by the lines $P_{prev}P$ and PD .

Then P deterministically forwards the message according to whether

$$\omega \geq \omega_{threshold}$$

where $\omega_{threshold}$ is some predefined parameter of the protocol.

In order to be able to implement the proposed algorithm, we assume that each sensor node has the ability to estimate the Direction of arrival (DoA) of incoming packets received from nearby nodes, and it can estimate the relative direction of the sink D , provided that the sink sends out beacon messages in regular intervals, to account for network mobility, and these can be captured by all the nodes in the network.

It is noted that the algorithm does not need any distance information (see the previous line of work [5] that analyzes a similar protocol under a stronger set of assumptions). Protocols requiring distance or exact location information can be implemented either with the use of GPS receivers that are unsuitable for sensor applications, by triangulation methods implied in neighboring nodes, which require severe bandwidth overhead, or by the deployment of acoustic sensors that add up to the overall node cost. Moreover, the proposed algorithm does not require nodes to have the same co-ordinates system. Such an approach would require the deployment of magnetic sensors (electronic compasses) that would again result in additional sensor cost. The algorithm is executed locally, and takes into account only local DoA measurements.

3 Experimental analysis

In order to analyze the performance of the algorithm described above, we performed a set of large scale experiments whose goal was to test the protocol's effectiveness under the following measures:

1. *Efficiency*: A small number (compared to the total number) of nodes must be activated for each data transmission towards the sink.




Fig. 1. (a) Ratio of activated nodes. (b) Success ratio

2. *Robustness or Fault tolerance*: The protocol should be able to deliver data to the destination, even when a large number of nodes is not responding (due to failure at deployment time)
3. *Scalability*: The routing algorithm should adapt easily to changes in the size of the network.
4. *Simplicity*: The algorithm must be able to run in an 8-bit microprocessor with minimal data memory (i.e. 4Kbytes).

It is obvious that the algorithm is extremely simple and since it does not depend on the network's size but only on its density (details omitted) it is also scalable so we will concentrate on proving its efficiency and robustness. We used the following setup for our experiments: $N = 2500$ sensors were spread uniformly at random in a square field, where all sensors have the same communication range R . For each simulation run we choose the sensor with the smallest and the largest $x - y$ coordinates to be the source and the sink, respectively. Hence, transmission will occur along the "diagonal" of the square field. Finally, to obtain valid statistical results all experiments were repeated a 1000 times.

Figure 1(a) shows the ratio r of activated sensors over the total number of sensors N as a function of the threshold angle, for $\omega_{threshold}$ equal to 90 up to 150 degrees. The six curves correspond to different communication ranges, each one resulting in a different average number μ of neighbors (8, 10, 12.5, 15, 17.5 and 20 neighbors, respectively).

As it can be seen in this figure, for small values of $\omega_{threshold}$ (around 90°) a large percentage of the sensors gets activated but the protocol is still better than flooding, especially for smaller values of μ . More importantly, however, the number of activated sensors decreases with the value of $\omega_{threshold}$ and becomes extremely small for suitable values of this quantity.

While the information shown on Figure 1(a) indicates that the protocol manages to activate only a few sensors for large values of $\omega_{threshold}$, it is important to know at the same time the success ratio of the protocol, i.e. the number of times the propagated information reached the sink. In particular, we would like to know whether there exists a value for $\omega_{threshold}$ where both the number of

activated particles is small and the success ratio is high. The answer to this question is shown on Figure 1(b).

It is obvious from this figure that the larger the average number of neighbors is the better the success probability becomes for any given threshold angle. Hence if we want to achieve a success ratio of 90%, we can either choose a threshold angle of about 95° when $\mu_{neighbors} = 8$, an angle of 115° when $\mu_{neighbors} = 10$, or an angle of 140° when $\mu_{neighbors} = 15$. So, one may ask: are all these settings equivalent? The answer of course depends on the number of activated particles. If we go back to Figure 1(a) and look at the corresponding curves, we will see that for these angles the number of activated particles becomes 50%, 40% and 25%, respectively. Hence using a forwarding angle of 140° in the $\mu_{neighbors} = 15$ setting, we manage to reach the destination 90% of the time while activating less than 25% of all the nodes.




Fig. 2. Size of propagating front of activated sensors

We have also been able to compute the average size of the propagating front of particles during transmission. The front is simply the particles that lie at the edge of the transmission zone towards the sink and is depicted in Figure 2. As can be seen from the figure the average front size decreases almost linearly with respect to $\omega_{threshold}$. When $\omega_{threshold} = 90^\circ$, the front size is less than $23R$, while at $\omega_{threshold} = 140^\circ$, the front size is less than $7R$ for all neighborhood distributions. Hence at this forwarding angle about 7 particles get activated on the average.

Finally, we investigated the fault-tolerance nature of our protocol when sensors die with various probabilities. The invariance of our algorithm under changes in the network size suggests the following approach: when we know that sensors may die with certain probability we can either plant more nodes or increase the communication range slightly to counteract the effect of dead nodes. In any case, using the results of Figure 1 we can optimize the algorithm's performance and obtain the required robustness (details omitted due to space restrictions).

4 Conclusions and Future Research

It is known from ancient times that the closest route between two points (at least in the Euclidean setting) is a straight line. In this work we have presented

SPEED, a Scalable Protocol for Efficient Event Delivery in sensor networks, where sensed data is sent to a receiving center using only local information and total absence of coordination. Our protocol is suited for those cases where unexpected changes to the environment (i.e. a fire, a person entering a restricted area, etc.) must be propagated quickly back to the base station without the use of complicated protocols that may deplete the energy of the network.

Our protocol is very easy to be implemented as it only requires nodes to decide whether or not to forward the message according to whether the angle formed between the previous node P_{prev} , the current node P and the destination D is greater than a predefined threshold angle $\omega_{threshold}$. We show that by carefully tuning this value one can obtain a tradeoff between the number of activated sensors and the success ratio of the protocol.

There are, however, some issues that need to be addressed. First, we need to overcome the problems related to the use of beacon frames for determining the relative direction of the sink as nodes in shadowed areas will be unable to receive the beacon frames. An initialization protocol is therefore required to ensure this information is shared by all nodes. Another issue is to increase the robustness of the protocol by assuring that data may reach the destination using more than one simple paths. This would make the protocol resistant to obstacles blocking communications between sensors. Finally, it is important to avoid depleting the energy of the sensors that lie in the optimal path between S and D (in the hypothetical scenario where S must send data to D frequently). Since the protocol is deterministic, the same sensors will be used continuously and hence they will be quickly drained of their energy. A solution to this problem is for the source S to choose a “random” path (not necessarily optimal) so that data is routed using this path. Of course, for this to be of any value, the protocol must again use only local information and no coordination between sensors.

References

1. I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, “Wireless sensor networks: a survey,” *Computer Networks*, vol. 38, pp. 393–422, March 2002.
2. Z. Haas, J. Halpern, and L. Li, “Gossip-based ad hoc routing,” in *IEEE InfoCom Proceedings 2002*, vol. 3, pp. 1707–1716, June 2002.
3. J. Kulik, W. R. Heinzelman, and H. Balakrishnan, “Negotiation-based protocols for disseminating information in wireless sensor networks,” *Wireless Networks*, vol. 8, pp. 169–185, March 2002.
4. C. Intanagonwiwat, R. Govindan, and D. Estrin, “Directed diffusion: a scalable and robust communication paradigm for sensor networks,” in *Proc. 6th International Conference on Mobile Computing and Networking*, pp. 56–67, 2000.
5. I. Chatzigiannakis, T. Dimitriou, M. Mavronicolas, S. Nikolettseas, and P. Spirakis, “A comparative study of protocols for efficient data propagation in smart dust networks,” in *International Conference on Parallel and Distributed Computing (EUROPAR)*, August 2002.
6. S. Tilak, N. B. Abu-Ghazaleh, and W. Heinzelman, “A taxonomy of wireless micro-sensor network models,” *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 6, pp. 28–36, April 2002.