

Efficient Event and Query Distribution in Sensor Networks

Man-Hon Chan, King-Shan Lui, Vincent Tam
Department of Electrical and Electronic Engineering
The University of Hong Kong
Pokfulam Road, Hong Kong, China
{mhchan, kslui, vtam}@eee.hku.hk

Abstract

A sensor network consists of a large number of sensors which are equipped with sensing, computation, and communication devices. Due to limitation in size, a sensor has only limited energy and storage. Traditional wireless network protocols cannot be applied in sensor networks directly. We study the distribution of events and queries in sensor networks. An event is something of interest detected by a sensor. A query is a request of information. A conventional approach to facilitate query nodes to acquire what they want is flooding. Nevertheless, flooding is not desirable in sensor networks due to the large number of nodes and limited energy in sensors. Recently, the concept of data-centric storage (DCS) is introduced where information of the same kind is kept in the same set of nodes. Queries can then be sent to these nodes for information retrievals. Theoretical analysis shows that this approach requires a lot fewer messages than the flooding approach when query frequencies are not high. Unfortunately, existing protocols based on DCS are prone to the hot-spot problem where some nodes have to handle lots of messages. In this paper, we present an efficient protocol for distributing events and queries in a location-aware sensor networks so that the load among nodes is more evenly distributed. We evaluate our protocol using simulations and the results show that our protocol successfully alleviates the hot-spot problem.

1 Introduction

Sensors are small devices that have capability to detect objects, collect information, and communicate with each other. A wireless sensor network usually consists of thousands of sensors which span a large geographical region. Sensor networks have become an important technology especially for environmental monitoring, military applications, disaster management, etc. Sensor networks differ from traditional wireless networks in many aspects and

many mechanisms developed for traditional wireless networks cannot be directly applied in sensor networks. In this paper, we study the problem of distributing events and queries in sensor networks. An event is something of interest, such as a sharp change in temperature, detected by a sensor. A query is a request of information of a certain event issued by a node in the sensor network. Different nodes may be interested in different events. A node knows what events it is interested in but does not know where these events happen. To get an answer, a query has to be sent to a node that possesses the information. A conventional approach to facilitate query nodes to acquire what they want is flooding. In the *push-based* approach, whenever an event happens, the information is flooded to all the nodes in the network. Then, every node knows the information when it needs it and does not need to send out any query at all. In the *pull-based* approach, a node that detects an event keeps the information in its own storage without sending the information to any other node. When a node wants that information, it floods the network with a query. The node that possesses the requested information replies the query.

Although flooding is a simple technique and is commonly used in many wireless network routing protocols, it is not very suitable for sensor networks. Sensor nodes are limited in energy and reducing the energy used in message transmission is one of the major optimization criteria of sensor network protocols. However, the number of messages required in each flooding is proportional to the number of nodes in the network. To make the problem worse, as the message will be broadcasted by nearby nodes at about the same time, collisions happen frequently and it leads to retransmissions of messages which consume extra energy. This problem is very serious in sensor networks since there are a large number of nodes and nodes are close to each other. As a result, flooding should be used only when necessary. Referring back to the *push-based* approach, every node receives the information of an event. However, it is not necessary if only a few number of nodes are interested in the information. Energy is wasted in distributing the event

information to nodes that do not require the information. The *pull-based* approach has a similar drawback. The query messages that are sent to nodes other than the node that detected the event are useless. It is obvious that energy can be saved if information is kept and distributed in a more intelligent manner.

To allow efficient routing of queries, the concept of *data-centric storage (DCS)* is introduced [11]. When a sensor detects an event, it decides which location is responsible for keeping the information and the information is sent to a node in that location. The location depends on the type of information and can be computed using a globally known hash function [9]. Then, when a node requests a particular piece of information, it can find out the location by the hash function and send a query to that location. A number of protocols have been developed based on the idea of DCS [9, 6, 4, 8, 10, 12, 3]. An important issue of DCS is where to keep event information since it affects how a query is propagated. This in turn affects the number of messages it takes to retrieve event information. For example, if an event is kept in a remote node in the network, a query may have to travel a long way, so does the answer message. On the other hand, if the event is kept in a node sitting in the center of the network, no matter which node issues a query, the query does not have to go through a very long path. Although keeping information in the center seems to reduce the message overhead of query propagation, keeping all the information in the center will highly increase the workload of the nodes there, since they have to handle all the queries. This hot-spot problem is a very serious problem in sensor networks since the sensors in the hot-spot will dissipate their energy very soon. As a result, it is necessary to strike a balance between total number of messages and the maximum number of messages a node has to handle.

In this paper, we describe a protocol for distributing event information and queries in location-aware sensor networks. We identify some problems of existing protocols and try to solve them in our protocol. We evaluate our protocol using extensive simulations. This paper is organized as follows: Related work is discussed in Section 2. Section 3 describes our network model and Section 4 presents our protocol. Simulation results are presented in Section 5 and we conclude our paper in Section 6.

2 Related Works

Some protocols have been developed based on the data-centric storage [9, 6, 8, 10, 12, 3].

[6] and [4] study how to efficiently answer queries about events that have scalar attributes. An example query provided in [6] is “list all events whose temperatures lie between 50° and 60° , and whose light levels lie between 10 and 15.” The authors develop a hash function such that

events having similar attribute values are kept in nearby nodes to allow more efficient querying.

[3] develops a *multi-resolution querying* scheme. An event is detected by a set of nodes instead of one. Different subsets of these nodes provide different resolutions of the event. Some nodes, called registration points, are designated to keep the sources of events but not the actual information of the events. When a node requests information about the event, it also has to specify the resolution it wants. The query is sent to the registration points and the registration points forward the query to the event sources. The event sources will then send replies to the requesting node based on the resolution requested.

The event model in this paper is similar to those used in [1, 7, 9, 8, 10, 12] that an event is something of interest and can be any kind of information. Events are independent of each other and occur randomly at anywhere.

Rumor Routing [1] and [7] do not follow the DCS approach. GHT [9] is one of the earliest works on event and query distribution for DCS. Each event is kept in a particular geographical location. This location is computed using a globally known hash function. Since it is possible that there is no node in the designated location, GHT develops a mechanism to assign a node in the neighborhood to keep the event. GPSR [5], a routing protocol that routes packets according to geographical locations instead of IP addresses, is enhanced to route messages to such a node. When a node searches for an event, it computes the location of keeping the event using the hash function and sends a query using the enhanced GPSR. GEM [8] follows similar idea except that it works in a network where geographical location information is not available. The authors propose to embed a tree in the network so as to provide relative location information. The authors also develop a routing algorithm to route messages on the tree.

It is not difficult to see that if the network is static and nodes never fail, GHT is able to answer all queries. However, when sensors are highly mobile, it is very likely that information is not kept in the calculated geographical location anymore by the time the query arrives. To solve the problem, [10] suggests to divide the network into geographical regions. Information of an event is kept in all nodes in that region instead of only one node. Then, as long as there is a node having the information stays in the region, a query can get an answer. Nevertheless, similar to [9] and [8], [10] has the hot spot problem. To simplify our discussion, we call the node that detects an event *event node*, the node that keeps the information *designated node*, and the node that issues a query *query node*. When an event is very popular, nodes surrounding the designated node and the designated node itself will have to transmit a lot of queries, which uses up energy quickly.

To provide different resiliency levels for different events,

keeping an event in multiple locations is suggested in [12]. More important information is kept in more hash locations to enhance resiliency. This mechanism alleviates the hot-spot problem to a certain extent due to the duplication of event storage. Nevertheless, an event may be sent to several locations which are very far apart, the number of messages to distribute an event is also a multiple of what are needed in GHT.

To solve the hot-spot problem, keeping the same information in multiple locations is the only way. To reduce the number of messages used to distribute the events in several locations, the locations should be related. In a location-aware sensor network, we distribute events only horizontally across the whole sensor network. Then, an event is stored in multiple sensor nodes but the number of messages needed is proportional to the dimension of the network. To find the requested information, a query is sent vertically to ensure it hits a node that contains the information. The next two sections describe our protocol in details.

3 Network Model

We focus on 2-dimensional sensor networks where sensors are lying on a rectangular region. We represent a sensor node as a point on the Cartesian plane. We assume that every node knows its geographical location before any event is detected or any query is generated. Each node keeps also the location information of its neighbors. We further assume that nodes know the boundary of the network and a node knows whether it is sitting on the perimeter or not. This information can be provided by a boundary detection mechanism such as [2]. We also develop our own perimeter node identification mechanism (Section 4.6).

We adopt the enhanced GPSR as the routing mechanism. Destination is specified as a geographical location. The original GPSR [5] is able to route a message to a node with exact geographical location as the destination. If there is no node with location exactly the same as the destination, the message will be dropped. The authors in [9] enhance GPSR to route the message to a node surrounding the destination that is closest to the destination among the nodes within the proximity. In our protocol, this enhanced GPSR is used for routing messages. For simplicity, we write GPSR, without “enhanced”, to refer enhanced GPSR when the context is clear.

There is a globally known hash function that hashes an event to a certain location, just as the one used in GHT. In GHT, an event is hashed to a 2-dimensional coordinate. In our protocol, we hash an event to an 1-dimensional coordinate. We call it *height level* since that is represented by y -coordinate as discussed later.

The location of a sensor node a is (x_a, y_a) . The hashed level of the event e is $y_{h(e)}$. The x -coordinates of the left

and right boundaries are x_l and x_r respectively.

4 Protocol

We present our protocol in detail in this section. As mentioned earlier, our protocol distributes events horizontally and queries vertically. The path that an event passes through is the event distribution path and the path that a query traverses is the query distribution path. When the event distribution path and the query distribution path intersect, which is the case in our protocol, the query can be answered. We now describe the event and query distribution of our protocol.

4.1 Event Distribution

To avoid the hot-spot problem, we would like the information of an event to be stored in a set of nodes instead of one node. On the other hand, these nodes should be spread out in the network instead of locating in a small region. To achieve this, our protocol distributes an event to the nodes which are on a similar height level across the whole network. That is, we first hash the event to a height level, say y' , and then distribute the information to the nodes with y -coordinate within a certain range of y' . The amplitude of the range depends on the topology and GPSR. We will discuss this amplitude issue further in Sections 4.3 and 4.4.




Figure 1. Illustration of Event Distribution

We now describe how an event can be distributed to nodes on a similar height level across the network and illustrate the procedure in Figure 1. When an event node a at location (x_a, y_a) wants to distribute an event e , it first finds out the height level of distribution, $y_{h(e)}$. a then sends the event using GPSR with destination $(x_a, y_{h(e)})$. Although there may not be a node locating exactly at $(x_a, y_{h(e)})$, GPSR will route the message to a node b , which is close to $(x_a, y_{h(e)})$ and is of “similar” height level as $y_{h(e)}$. b then distributes the message on that height level. It achieves this

by sending a copy of the message to $(x_l, y_{h(e)})$ and another copy to $(x_r, y_{h(e)})$. Only the nodes on a similar height level as b have to store the event in memory. That is, only the nodes on the paths from b to c and from b to d have to keep the information. Nodes of the path from a to b do not need to keep information.

There may not be nodes sitting exactly at locations $(x_l, y_{h(e)})$ and $(x_r, y_{h(e)})$, such as the situation in Figure 1. Although GPSR is able to stop searching after finding the closest node to the destination, our simulations show that GPSR may go around the whole network before it terminates at such a node. To avoid this redundant searching, we require a perimeter node to stop propagating an event even though it is not sitting exactly at the destination specified in the packet. For example, node c in Figure 1 will stop propagating the event message since it is a perimeter node. The same principle applies to node d .

There are some minor details about event distribution and those will be described in Sections 4.3, 4.4 and 4.5.

4.2 Query Distribution

We now describe the query distribution. When a node wants to find the information of the event e . Using the hash function, it knows that nodes sitting on a level close to $y_{h(e)}$ can provide an answer. Let the query node q be at location (x_q, y_q) . To request the information, q sends a query to $(x_q, y_{h(e)})$. When a node receives the query, it checks whether it has the requested information. If yes, it should reply the query; otherwise, it should continue to forward the query.

As mentioned in Section 3, GPSR will route a message to a node that is closest to the destination location. It is not guaranteed that the closest node to $(x_q, y_{h(e)})$ contains the information of event e . It is illustrated in Figure 2. The white nodes are the nodes that contain the information. They are the nodes on the event distribution path in Figure 1. When q sends a query to $(x_q, y_{h(e)})$, GPSR routes it to p . Nevertheless, no node on path q to p has the requested information.

Fortunately, some neighbors of p have the information of event e . p can then obtain the information by sending a broadcast message to its neighbors. A neighbor, say c , replies p with the information and p can forward it back to q . If p does not hear anything from its neighbors, p replies *no-such-event* to q after timeout.

Although the broadcasting performed by p stops at its neighbors, this is still not desirable since it is possible that all neighbors send a reply and it is a waste of energy. To reduce the need of broadcasting, we adopt the *promiscuous mode* in event distribution, which will be described in the next section.




Figure 2. Illustration of Query Distribution

4.3 Promiscuous mode

When a node sends an event message to its neighbor according to GPSR, nodes in the proximity can overhear the message. In the promiscuous mode, nodes overhearing a message also register the event in their memories [1]. By using promiscuous mode, the need of broadcasting as mentioned earlier can be reduced.

To see this, let's consider the situation in Figure 3. Since the white node does not have the event, the query is further forwarded upwards. If the white node saves the event information when it overhears the event message, it can answer the query right away.




Figure 3. Advantage of Promiscuous Mode

In essence, we convert the event spreading path into an event spreading strip. This not only shortens the distribution path of queries, but also helps to ease the hot-spot problem. Thus, more nodes in the area can answer the same query.

However, promiscuous mode has some drawbacks. Each node would carry too many events which probably exceed their memory limits. To strike a balance, in our protocol, a node a keeps the event e it overhears only when $|y_{h(e)} - y_a| < \frac{range}{3}$, where *range* is the maximum overhearing distance. This method effectively narrows the strip and reduces redundancy.

4.4 Event Distribution Path Straightening

Our query distribution for event e ends at a node sitting on the height level around $y_{h(e)}$. Whether the query gets an answer depends on whether there is a node around level $y_{h(e)}$ possesses the answer. If the event distribution path is a straight line across level $y_{h(e)}$, there should be no problem in getting a reply for any query. Unfortunately, it is intrinsically impossible to spread an event message in a straight horizontal line. Our simulations show that GPSR may not always forward the event message along a straight line and it affects our query hit ratios.

To solve the problem, we apply a straight line heuristic to guide the event distribution to stay on the line $y = y_{h(e)}$ as much as possible. Suppose that the destination d of a message is at location (x_d, y_d) . To identify a neighbor to be the next hop for forwarding the message, GPSR calculates for each neighbor n at (x_n, y_n) the distance between n and d , which is $\sqrt{(x_d - x_n)^2 + (y_d - y_n)^2}$. GPSR selects a neighbor with the smallest distance to d as the next hop. To form a straight path, we should not select any next hop neighbor n such that y_n differs a lot from y_d . To do that, we modify the distance calculation in GPSR to

$$\sqrt{(x_d - x_n)^2 + 9(y_d - y_n)^2}$$

that essentially exaggerates the deviation in y . Since GPSR is greedy, it favors points nearer to our desired contour. As a result, the event spreading path will be much straighter than using the original distance calculation mechanism.

Rumor Routing [1] has also introduced another straightening algorithm where the event message would recognize recently seen neighbors. In search of the next hop, it chooses a node not on the list. We claim that our approach is better than this algorithm because we take real geographical information into account and do not have to construct large packets to store recently seen nodes.

4.5 Random Biasing

GPSR performs poorly when the destination is outside the perimeter. Suppose that a node a detects an event e . To distribute the event, node a first sends the event message to $(x_a, y_{h(e)})$. If a is at the boundary, $(x_a, y_{h(e)})$ is likely to be outside the network. GPSR performs badly in such case where it must route through the entire perimeter.

Our primary goal is to route the message to a node on a level around $y_{h(e)}$, and the x -coordinate of this node is not significant. To solve this problem, we choose to bias our destination x -coordinate towards the center by a random factor. This random factor is calculated by multiplying the x -coordinate distance between a and the center and a random number, which is arbitrarily set to be at most 0.2. This is based on two reasons:

1. There are likely more nodes in the central region. Although we route messages to the center area, the extra load can be easily distributed to all nodes, and does not add much to the overall hot spot.
2. The random factors help distribute the workload and hot spot over a larger set of nodes. Even if a single node is launching queries frequently, the query would be distributed to different regions, and no specific region would be overloaded.

4.6 Perimeter Detection

In our algorithm, the event spread halts on the perimeter. In the original GPSR, event message destined to a location on the boundary where no node exists would result in routing through the perimeter of the entire network. We avoid this by deliberately stopping the spread of an event message when the event message hits a perimeter node. While perimeter nodes can be found by some existing algorithms, we develop a simple mechanism based on GPSR that works with our protocol. The steps of the mechanism are as follows:

1. An arbitrary node starts the mechanism by routing a message to $(x_l - 1, 0)$. This point is beyond the bound of the network and is thus routed around the perimeter by GPSR. GPSR halts when it visits the same edge twice. After it halts, we know that the packet has been routed around all perimeter nodes of the *outer face* once. Let the node where GPSR terminates be p .
2. p sends a packet to $(x_l - 1, 0)$. GPSR would route the packet via the same route as above in this step. We mark every sensor node along the path as perimeter. In this way, all nodes on the perimeter of the *outer face* are then marked as on the perimeter.

5 Simulation

We evaluate the performance of our protocol and compare it with GHT using simulations.

5.1 Constrained hashing

The hash function used is not well defined in the GHT paper. We have tried to employ a random uniform hashing that spans across the whole network. We generate $x_{h(e)}$ such that it is uniformly distributed over $[x_l, x_r]$ and $y_{h(e)}$ that is uniformly distributed over $[y_t, y_b]$, where y_t and y_b define the upper and lower bound of the network area. Essentially, we utilize the full range of our network boundary. However, the result of GHT turn out to be relatively poor

because a significant portion of the hashed coordinates reside outside the network. To fix the problem, we employ a constrained version of hashing in our simulations. $x_{h(e)}$ is distributed over $[x_l + width * 5\%, x_r - width * 5\%]$ and $y_{h(e)}$ is distributed over $[y_t + height * 5\%, y_b - height * 5\%]$. We basically avoid nodes sitting too close to the boundary to serve as data storage.

5.2 Simulation Setup

The simulation setup is summarized in Table 1. We generate 10 networks for each topology setup. That is, we generate 10 different networks where the number of nodes is 100 and the average number of neighbors of the nodes is between 5 and 8. We simulate different event to query ratios for the 80 networks generated. There are 50 events and they all happen in random locations. We try launching different numbers of queries. Queries are generated randomly, both in locations and the events they are asking for. That is, different events may have different numbers of nodes requesting information.

Number of Nodes	100, 200, 500, 1000
Average number of neighbors	5-8, 8-12
Number of Events	50
Number of Queries	50, 250, 500

Table 1. Simulation Setup

5.3 Simulation Results

We measure (1) the number of messages involved in the event, query and answer distributions (NM), and (2) the maximum number of messages a node has to handle (HS). The first one reflects whether the protocol is energy efficient while the second one reflects the seriousness of the hot-spot problem. The results are shown in Tables 2, 3, and 4.

Size	NM(Our)	NM(GHT)	HS(Our)	HS(GHT)
100	1745	2184	46.3	55
200	1800	2049	27.4	40.6
500	2300	2211	39.8	53.4
1000	3258	2921	61.7	83.4

Table 2. 50 Events and 50 Queries

It can be observed that our protocol requires fewer number of messages while reducing the maximum number of messages a node has to handle in most cases. We therefore conclude that our protocol is an efficient protocol for event and query distribution in sensor networks.

Size	NM(Our)	NM(GHT)	HS(Our)	HS(GHT)
100	3054	3500	96.4	93.4
200	3343	3758	73.4	103.8
500	4658	4999	66.0	133.0
1000	6907	7296	133.9	208.3

Table 3. 50 Events and 250 Queries

Size	NM(Our)	NM(GHT)	HS(Our)	HS(GHT)
100	4653	5119	160.2	153.9
200	5291	6649	117.1	177.8
500	7635	9756	102.3	208.0
1000	11469	14580	188.8	316.2

Table 4. 50 Events and 500 Queries

6 Conclusion

In this paper, we study the problem of distributing events and queries in sensor networks using data-centric storage. We firstly identify the weaknesses of existing approaches and then describe our new protocol which distributes events horizontally and queries vertically. We measure the performance of our protocol using extensive simulations and compare with GHT, a pioneer work in this area. Simulation results show that our protocol outperforms GHT in most situations. Therefore, we conclude that our protocol is promising for data-centric storage in sensor networks.

References

- [1] D. Braginsky and D. Estrin. Rumor Routing Algorithm For Sensor Networks. In *ACM International Workshop on Wireless Sensor Networks and Applications (WSNA)*, pages 22 – 31, 2002.
- [2] B. Carburnar, A. Grama, and J. Vitek. Distributed and dynamic voronoi overlays for coverage detection and distributed hash tables in ad-hoc networks. In *IEEE International Conference on Parallel and Distributed Systems (ICPADS)*, 2004.
- [3] J. Chen, Y. Guan, and U. Pooch. An efficient data dissemination method in wireless sensor networks. In *IEEE Global Telecommunications Conference (Globecom)*, volume 5, pages 3200 – 3204, 2004.
- [4] B. Greenstein, D. Estrin, R. Govindan, S. Ratnasamy, and S. Shenker. DIFS: A Distributed Index for Features in Sensor Networks. In *IEEE International Workshop on Sensor Network Protocols and Applications (SNPA)*, 2003.
- [5] B. Karp and H. T. Kung. GPSR: Greedy Perimeter Stateless Routing for Wireless Networks. In *ACM Mobicom*, pages 243 – 254, 2000.
- [6] X. Li, Y. J. Kim, R. Govindan, and W. Hong. Multi-dimensional Range Queries in Sensor Networks. In *ACM SenSys*, pages 63 – 75, 2003.

- [7] X. Liu and Q. H. Y. Zhang. Combs, Needles, Haystacks: Balancing Push and Pull for Discovery in Large-Scale Sensor Networks. In *ACM SenSys*, 2004.
- [8] J. Newsome and D. Song. GEM: Graph EMbedding for Routing and Data-Centric Storage in Sensor Networks Without Geographic Information. In *ACM SenSys*, pages 76 – 88, 2003.
- [9] S. Ratnasamy, B. Karp, L. Yin, F. Yu, D. Estrin, R. Govindan, and S. Shenker. GHT: A Geographic Hash Table for Data-Centric Storage. In *ACM International Workshop on Wireless Sensor Networks and Applications (WSNA)*, pages 78 – 87, 2002.
- [10] K. Seada and A. Helmy. Rendezvous Regions: A Scalable Architecture for Service Location and Data-Centric Storage in Large-Scale Wireless Networks. In *International Parallel and Distributed Processing Symposium (IPDPS)*, 2004.
- [11] S. Shenker, S. Ratnasamy, B. Karp, R. Govindan, and D. Estrin. Data-Centric Storage in Sensornets. In *ACM SIGCOMM HotNets*, 2002.
- [12] R. Tamishetty, L. Ngoh, and P. Keng. An efficient resiliency scheme for data centric storage in wireless sensor networks. In *IEEE 60th Vehicular Technology Conference*, volume 4, pages 2936 – 2940, Fall 2004.