Poster: Sink-to-Sensors Reliability in Sensor Networks*

Seung-Jong Park Raghupathy Sivakumar
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332, USA
sjpark@ece.gatech.edu, siva@ece.gatech.edu

ABSTRACT

The issue of reliability, thus far, has not been addressed thoroughly to any extent in sensor networking research. To the best of our knowledge the first such work that focused on some aspect of reliable communication in sensor networks appeared only in [1]. In this paper, we focus on the problem of communication reliability from the sink to the sensors in a static network.

1. WHY RELIABILITY?

The need (or lack thereof) for reliability in a sensor network is firmly dependent upon the specific application the sensor network is used for. Consider a sensor network deployed to detect the presence of harmful gases in an occupied building, with the sink having the capability of issuing queries indicating what specific gas the sensors should attempt to detect. Given the nature of the application, it is absolutely critical that a query update reach the sensors in a reliable manner. As we elaborate later, the specific form of reliability might change from application to application. However, any sensor network that is deployed to cater to a critical application (in both civilian and military environments) will require mechanisms to ensure reliable delivery of information from the sink to the sensors. Besides delivery of queries, reliability will also be required when control software (e.g. a new data processing tool or a new operating system) is downloaded to upgrade the sensors.

2. WHAT TYPES OF RELIABILITY?

Different types of reliable delivery can be required in a sensor network depending upon the nature of the message, and the scope of the delivery. We distinguish the first two types of reliability based on the size of the message: (i) **Single packet message delivery:** Most queries in a sensor network can be expected to be small enough to fit on a single packet. Ensuring single packet delivery is significantly more challenging than ensuring multiple packets delivery as receiver initiated NACK schemes cannot be employed, since receivers do not know that a message has been transmitted to be reliably received by them. Hence approaches such as those proposed in [1] will fail when a single packet has to be reliably delivered to the sensors. (ii) **Multiple packets message delivery:** This corresponds to either multiple packet queries or larger messages such as control software deliveries. Such message deliveries are relatively more tractable to achieve due to the inherent redundancy present in the form of multiple packets. As long as a sensor receives one part of the message, it can actively participate in fetching the other parts of the message.

The scope of the message delivery can be used to classify the types of reliability further: (iii) **Message delivery to all sensors:** This is most straightforward type of delivery, where the message needs to be delivered to all sensors in the network. (iv) **Message delivery to all sensors in a sub-area.** It is possible that the sink wants to deliver a message to only sensors in a particular sub-area. We restrict the focus of this work to sub-areas that are contiguous regions including the sink. (v) **Message delivery to *cover* entire sensor area or sub-area.** This is different from the delivery to all sensors due to the typical redundant deployment of sensors. Essentially, the message from the sink needs to be reliably delivered such that sensors that receive the messages, among themselves, cover the entire sensing field.

3. WHAT CHARACTERISTICS NEED TO BE CONSIDERED?

The characteristics of sensor networks that need to be accounted for in developing a solution to the problem of reliable message delivery are: (i) **Scale:** The scale of a sensor network is several orders of magnitude greater than that of typical ad-hoc network environments. Hence, the issue of not using global coordination in any form is highly critical. (ii) **Density:** In sensor environments, where the node density can be expected to be very high, the broadcast storm problem will be severe, requiring robust reliability mechanisms. (iii) **Scarcity of resources:** Sensor networks, much like their ad-hoc counterparts, are characterized by scarce bandwidth, and battery power. Hence, the reliability mechanisms will have to be highly efficient. (iv) **Timeliness:** A query that needs to be reliably delivered is also more likely to have strict constraints in terms of the tolerable delivery latency.

4. WHAT ARE KEY DESIGN DECISIONS?

We now present key design decisions that are motivated by the characteristics of sensor networks: (i) **ACK vs. NACK:** Using an ACK based reliability scheme will result in the well known ACK implosion problem, which can especially be severe in sensor networks due to its scale. Hence, a NACK scheme is preferable. How-

*This work was funded in part by NSF grants ANI-0117840 and ECS-0225497, Yamacraw, and the Georgia Tech Broad-band Institute

MobiHoc'03, June 1–3, 2003, Annapolis, Maryland, USA.

ever, a NACK based scheme still needs to address the issue of a possible NACK implosion in a local neighborhood, and has to be supplemented with other mechanisms to provide for single packet reliable delivery. (ii) **Local vs. Non-local recovery:** Given the typical size of a sensor network, local recovery is by far more preferable. (iii) **Designated vs. non-designated recovery servers:** The advantage of designated servers is the potential performance improvements that can be gained by appropriate designation. A non-designated server approach, on the other hand, can result in severe overheads due to the lack of any coordination. (iv) **Dynamic designation vs. static designation:** Due to the dynamic nature of sensor networks, it is essential that the designation of recovery servers are done as close (in time) to the message transmission as possible. (v) **Out-of-sequence vs. in-sequence propagation:** A NACK based scheme coupled with propagation of out-of-order packets belonging to a message can trigger unnecessary NACK messages being transmitted by all downstream sensors. On the other hand, propagating only in-sequence packets (like in [1]) can result in the wastage of precious bandwidth resources downstream.

5. WHAT IS THE IDEAL APPROACH?

We briefly present an ideal approach for designating recovery servers that assumes full knowledge of the loss pattern for a single packet flood. Ideally, the recovery servers should be chosen on a per-packet basis, although such an approach is evidently undesirable from a feasibility standpoint. The problem of ideal server designation, given the loss pattern, can be solved as follows. Consider all nodes that have received the packet to be colored black, and all nodes that have not received the packet to be colored white. For simplicity, consider a scenario where there are no islands of white nodes (in other words, every white node has at least one black node in its neighborhood). The ideal server designation solution is then to choose the minimum number of black nodes that will cover all the white nodes in the network. This is a variation of the **set cover problem** that has been shown to be NP-hard [2]. While sophisticated approximations are possible, a simple greedy algorithm would involve choosing a black node with the maximum white-degree, changing the color of the white neighbors of that node to black, and repeating the process till there are no white nodes (note that this process would handle white-islands also). If only the chosen nodes (re)transmit the packet, the number of retransmissions would be minimized. A desirable side-effect of such an approach is the minimization of the probability of such retransmissions colliding with each other. The design decisions outlined in Section 4 and the the ideal recovery server designation approach described above together produce an ideal solution for achieving reliable sink-to-sensors delivery.

6. HOW CAN THE IDEAL SOLUTION BE APPROXIMATED?

We propose a two-pronged distributed solution for achieving reliable delivery of messages in a sensor network:

(i) For single packet messages, we propose a two-radio solution where each node is equipped with a low frequency ”busy-tone” radio (R2) in addition to the default radio (R1) that is used for data transmissions and receptions. The busy-tone radio can only be in one of two states - ON or OFF. When the sink initiates a message, it first switches ON its R2. When any sensor in the network hears an R2 signal, it switches ON its R2. A sensor turns OFF its R2 only upon receiving the first packet of the message. Since the state of the R2 channel is purely binary in nature (ON or OFF), it does not experience the typical problems of collisions and contention.

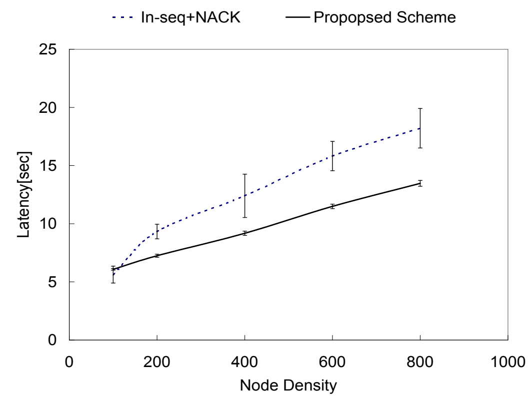

Figure 1: Latency for all sensors to receive all data, for 20 scenarios, with 95% confidence interval

The presence of an R2 signal acts as an implicit NACK, triggering retransmissions from nodes that hear the signal. The same mechanism is used even for a a multiple packets message, although only for the first packet. The remaining packets are reliably delivered by the second part of our solution.

(ii) For reliable delivery of all packets except for the first packet, we propose an approximation to the solution for the Set Cover problem that involves the computation of the minimum dominating set (MDS) of the underlying network, independent of the loss occurrences of any particular packet. In keeping with the goal of addressing the dynamic nature of sensor networks, the minimum dominating set is computed afresh for every new message that is transmitted by the sink, and needs to reliably delivered. The fact that the first packet of a message will be reliably delivered through the two-radio scheme is leveraged in the computation of the MDS. Unlike any other existing MDS computation scheme, the proposed algorithm will compute the MDS in the course of a single packet flood. Once the MDS constructed, loss recovery first occurs for the MDS nodes. Since packets are forwarded out-of-order also, the MDS nodes exchange meta-information in the form of bit-maps to prevent unnecessary retransmission requests. Once the MDS nodes recover from all losses, they retransmit the packets lost by any of their non-MDS neighbors. It can be shown that such an approach incurs only a constant order times more overhead than the ideal solution where the server designation is done on a per packet basis.

7. HOW DOES THE PROPOSED APPROXIMATION PERFORM?

Figure 1 shows a representative result that compares the approach presented in [1] with that of the proposed scheme, for different network node densities. It can be observed that the time-delay incurred in successfully delivering the entire message reduces by up to 40% for the proposed approach.

8. REFERENCES

[1] C-Y. Wan, A. Campbell, and L. Krishnamurthy, “PSFQ: A Reliable Transport Protocol for Wireless Sensor Networks,” in *Proc. ACM Int. Workshop on Sensor Networks and Architectures*, Atlanta, Sept. 2002.

[2] David S. Johnson, “Approximation algorithms for combinatorial problems,” in *J. Computer and System Sciences 9*, 1974, pp. 256–278.