Scalable Landmark Flooding - A Scalable Routing Protocol for WSNs

Pengfei Di and Thomas Fuhrmann
Computer Science Department, Technische Universität München, Germany
di@so.in.tum.de, fuhrmann@so.in.tum.de

ABSTRACT

Wireless sensor networks (WSNs) are about to become a popular and inexpensive tool for all kinds of applications. More advanced applications also need end-to-end routing, which goes beyond the simple data dissemination and collection mechanisms of early WSNs. The special properties of WSNs – scarce memory, CPU, and energy resources – make this a challenge. The Dynamic Address Routing protocol (DART) could be a good candidate for WSN routing, if it were not so prone to link outages.

In this paper, we propose Scalable Landmark Flooding (SLF), a new routing protocol for large WSNs. It combines ideas from landmark routing, flooding, and dynamic address routing. SLF is robust against link and node outages, requires only little routing state, and generates low maintenance traffic overhead.

Categories and Subject Descriptors: C.2.2 [Network Protocols]: Routing protocols

General Terms: Algorithms.

Keywords: Ad-hoc routing, Landmark, Flooding.

1. INTRODUCTION

DART [1] is a simple routing protocol for ad-hoc and mesh networks. Its core idea is to assign the nodes their addresses according to the physical network topology. A node joining the network obtains its address from its new neighbors. The address assignment is such that nodes with common address prefixes form connected areas. Routing tables then only need to contain routes to log n other nodes: The node with address 1011 has routes to one node in 0xxx, 11xx, 100x, and to 1010 – if that nodes exist. These nodes serve as relays for their respective area.

The authors of DART propose to use a distance vector routing protocol (DVR) to obtain the routes. But DVR has well-known problems with resilience in case of link failures. Link state routing protocols are more robust, but create more route maintenance traffic. Moreover, DART assumes that the network builds up gradually rather than being switched on as it is often the case with WSNs.

In order to circumvent the shortcomings of DART, we propose to use flooding to assign the addresses and populate the routing tables. We exploit DART's hierarchical segmentation of the network into areas and show that it suffices when only one node floods each area. As a result, the control message overhead is only $O(\log N)$ rather than $O(N)$ as it would be with ordinary flooding.

CoNEXT Student Workshop'09, December 1, 2009, Rome, Italy.
ACM 978-1-60558-751-6/09/12.

2. SCALABLE LANDMARK FLOODING

Like DART, SLF segments the WSN into a binary tree hierarchy of areas. In the example of fig. 1, the entire network area A is divided into two sub-areas A_0 and A_1, A_0 is sub-divided into A_{00} and A_{01}, etc. In each area one of the nodes serves as landmark for this area. Its address is assigned such that it corresponds to the area ID filled-up with zeros. In our example, node N_{000} serves as landmark for A_0 and A_{00}. Node N_{010} serves as landmark for A_{01}, and N_{100} serves as landmark for A_1.

This segmentation is a recursive process that is triggered by the initial deployment of the network. By default, the WSN's sink node is assigned the address N_{000}. Upon bootstrapping, it floods the network with a *landmark announcement* (LA) message.

Segmenting an Area

When receiving an LA message, each node enters the following information into its routing table (cf. table 1):

- the area covered by this LA message (here A_0),
- the source address from where the LA message was emitted (N_{000} in our case),
- the hardware address of previous forwarder, and
- the hop count of the message.

Due to the broadcast nature of the WSN links, a node will receive multiple copies of each LA message. This duplicate information is not entered into the routing table, unless it contains a lower hop count. In that case the new information overwrites the previous entry. The duplicate LA messages are, however, used to build up a table of direct neighbors.

A time T_{wait} after having flooded the network, the source node – N_{000} in our example – emits a *landmark solicitation* (LS) message. The nodes forward this message in a random-walk manner downstream away from its source. If a node

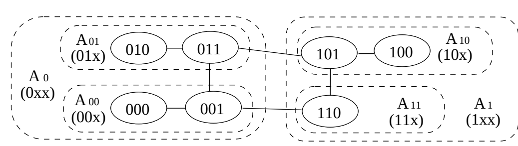

Figure 1: Address and area assignment

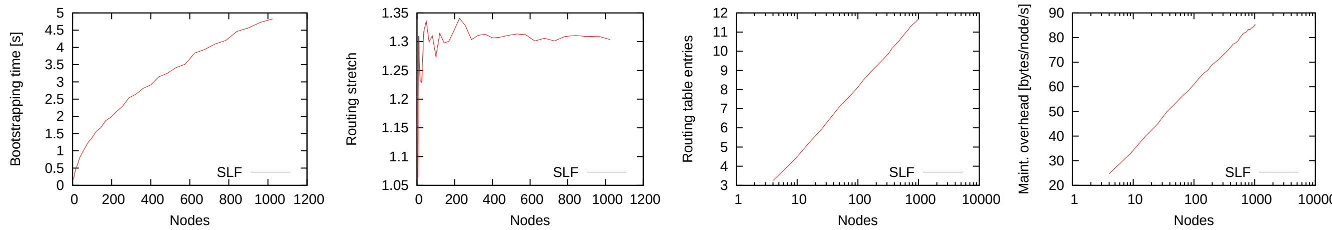

Figure 2: First performance evaluation of SLF (lattice topology, 10ms link delay and 3s maintenance interval)

has no neighbor that is farther from the source than itself, the node becomes the antipode landmark of the soliciting node.

In our example, N_{000} flooded A to announce itself as landmark for A_0. Its antipode becomes the landmark for A_1 and accordingly assumes the address N_{100}.

The antipode landmark then floods the network with an *antipode landmark announcement* (ALA) message, so that each node in A knows its distance to N_{000} and N_{100}. Nodes closer to N_{000} become members of A_0, and vice versa. Ties are broken as follows: If the landmark and its antipode are reached via the same neighbor, a node must adhere to that node's decision. Otherwise it may choose randomly one of the two landmarks. This guarantees that the area is divided into two connected sub-areas. Since this decision can be made as soon as a node receives the antipode's LA message, a node's decision can be piggy-backed onto the ALA message.

Recursively Segmenting the Network

After having received the ALA message, N_{000} emits another LS message, this time restricted to A_0. Accordingly, T_{wait} after having flooded the network with its ALA message, N_{100} emits an LS message that is restricted to A_1.

Within the sub-areas, the same method as above applies. After $O(\log N)$ rounds, the sub-areas contain single nodes only, and the address assignment ends.

Ideally, each round would split the network in half, and the process would end after exactly $\log N$ rounds. But the random walk can of course not guarantee to find the optimal antipodes. In the worst case, a star with the sink as hub, the area hierarchy would degenerate.

The set-up phase requires $2 + N^{-\frac{1}{2}}$ messages per node and round: an LA and an ALA message that are flooded, and an LS message that travels about $\sqrt{N}$ hops during its random walk. Thus, altogether, we have $O(\log N)$ messages per node for the set-up.

Routing

As in Pastry and in DART, routing proceeds from the most to the least significant bit of the address. Packets in A_0 that are destined to a node in A_1 are routed towards the landmark N_{100} in A_1. Or, to put it differently, a node N_{0xx} forwards packets with destination N_{1xx} towards N_{100}.

Sub-area	Landmark	Next-hop	Hop Count
A_1 (1xx)	100	MAC_{101}	2
A_{00} (00x)	000	MAC_{001}	2
A_{010} (010)	010	MAC_{010}	1

Table 1: Routing table at node 011

Maintenance and Stability

Landmarks and antipode landmarks regularly flood their super-area with *landmark maintenance* (LM) messages: N_{000} and N_{100} flood A, N_{010} floods A_0, N_{110} floods A_1, etc. Nodes use these messages to update their routing tables.

Assume that in fig. 1 the link between N_{001} and N_{110} breaks. From the flooding, N_{001} learns that N_{011} is the next hop towards A_1 and updates its routing table accordingly. Such a route update is fast and inexpensive.

If, however, the link between N_{001} and N_{011} breaks, it does not suffice to update the routing tables, because, e. g. , N_{001} cannot reach N_{011} within A_0 any more. In such a case, the nodes must reconfigure the addresses. The nodes can detect this case from missing LM messages. In our example, N_{000} would not hear the LM messages from N_{010} any more. Thus, after a timeout, it will flood the network with an LA message to trigger a new address initialization process.

In practice, we assume a WSN to be so well connected that this kind of network split on the top level is rare. If it occurred on a lower level, the re-keying would be limited to that area only.

The LM messages also allow the nodes to detect the loss of a landmark: If N_{000} does not hear LM messages from N_{100} or N_{010} any more, it triggers an address initialization process. The same applies to the lower levels of the area hierarchy. Only the absence of N_{000} cannot be detected in that way. If we want to protect against a failing sink node, too, we can extend the protocol with a mechanism to elect another N_{000} landmark.

3. CONCLUSION AND FUTURE WORK

Scalable Landmark Flooding is a scalable end-to-end routing algorithm for large WSNs. It has a low complexity of per-node state and maintenance overhead. More importantly, since the routing information is always up-to-date, we expect SLF to be more robust against node and link failures than many other ad-hoc routing protocols.

As a first evaluation, we have simulated SLF in the OMNeT++ simulator using a lattice topology (cf. fig. 2). Currently, we evaluate it in comparison to DART, beacon vector routing, virtual ring routing, and scalable source routing. This evaluation shall also explore if we can efficiently apply DART's address look-up mechanism in our context.

4. REFERENCES

[1] J. Eriksson, M. Faloutsos, and S. V. Krishnamurthy. DART: Dynamic address routing for scalable ad hoc and mesh networks. *IEEE/ACM Transactions on Networking*, 15:119–132, 2007.