PL-Gossip: Area Hierarchy Maintenance in Large-Scale Wireless Sensor Networks

[Invited Poster – Extended Abstract]*

Konrad Iwanicki (student)†‡ and Maarten van Steen (faculty)†

†Vrije Universiteit, Amsterdam, The Netherlands

‡Development Laboratories (DevLab), Eindhoven, The Netherlands

{iwanicki, steen}@few.vu.nl

http://www.few.vu.nl/~iwanicki/lupa-web/

1 INTRODUCTION

Many existing and novel wireless sensor network (WSN) applications require large numbers of sensor nodes operating over long periods of time. The effort involved in the deployment and durable maintenance of such networks can be reduced if nodes autonomously organize themselves into a required logical network structure.

A typical way of self-organizing the network in the first-generation WSN applications was having nodes build and maintain a tree rooted at a base station. To ensure scalability, however, a rapidly growing number of compelling WSN applications employs a different network structure: a recursive geometric organization. Examples of such applications include in-network storage [2], reactive tasking [3] (based on local observations, sensor nodes trigger actuator nodes), object tracking [4], network health monitoring [5], and various query engines, like multi-dimensional range queries [6], spatial range queries [7], or multi-resolution queries [8].

Recursive geometric network organization exploits node proximity and connectivity. It provides scalable recursive naming of network areas, that is, we can name a network area, the subareas of this area, and so forth. Moreover, the structure enables routing between any of such areas or between any pair of nodes.

Devising a protocol in which nodes autonomously build and maintain such an organization poses a number of challenges. The combination of a possibly large network size and a very short radio range of sensor nodes leads to high-diameter multi-hop topologies. Because of the memory and bandwidth limitations of individual nodes, the state stored by each of them must scale gracefully with the network size. To enable predicting the network lifetime and provisioning and conserving the battery power accordingly, the protocol must ensure predictable or well-defined maintenance traffic. Yet, the maintenance must guarantee adaptability to continuous network dynamics. For practical reasons, nodes should be able to build and maintain the organization in many heterogeneous settings, ranging from “planar” regions (e.g., parking lots), to “volumetric” deployments (e.g., interiors of multi-story buildings). To the best of our knowledge, none of the existing solutions meets all of the above goals.

2 OUR APPROACH

Our protocol, dubbed PL-Gossip [1], enables nodes to self-organize into a practical instance of a recursive geometric network structure, known as an area hierarchy [9]. The essential idea is that the nodes autonomously group themselves into sets based on their connectivity. The groups correspond to network areas and form a multi-level hierarchy (see Fig. 1) that provides an addressing scheme and enables efficient routing and multicasting. The membership of a node in the hierarchy is reflected by the label of this node, which also constitutes the node’s routing address. This scheme requires only $O(lgN)$ state per node and works equally well in both planar and volumetric multi-hop deployments.

To ensure predictable traffic and at the same time to enable adapting the group hierarchy to changes in the node population and connectivity, the protocol em-

*Further information can be found in an accompanying technical report [1], available online.

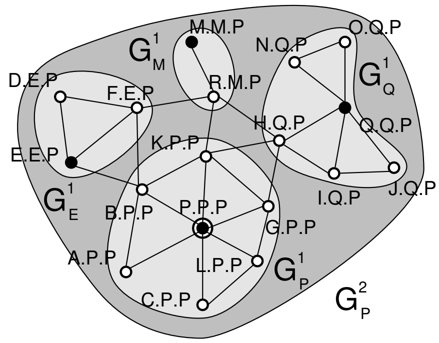

Fig. 1: *An example of a group hierarchy with the corresponding nodes' labels.*

ploys periodic gossiping. From a node's perspective, the time is divided into rounds. In every round, each node broadcasts a single beacon message and receives similar messages from neighboring nodes. A beacon contains the label and the routing table of the sender (plus some necessary consistency information). Such a simple, well-defined traffic pattern, which is a distinguishing feature of PL-Gossip, provides two important properties. First, it allows for precise scheduling of radio activity periods by the nodes. This is crucial for power provisioning as a sensor node with an active radio consumes several orders of magnitude more battery power than a node having its radio off. Second, gossiping makes the protocol extremely robust against massive node failures, network partitions, etc. The group hierarchy is restored even when a number of such events occurs concurrently.

This gossip-based traffic pattern, however, requires the nodes to build and maintain the group hierarchy using only their local information and the information received in the beacon messages of their neighbors (in particular, no messages are forwarded). PL-Gossip copes with this stringent limitation by defining special constraints on the hierarchy, as formalized below.

Property 1. *Level* 0 *groups correspond to individual nodes.*
Property 2. *There exists a single, level* $\mathcal{H}$ *group that contains all nodes.*
Property 3. *Level* $i+1$ *groups (where* $0 \leq i < \mathcal{H}$*) are composed out of level i groups, such that each level i group is in exactly one level* $i+1$ *group.*
Property 4. *Each level* $i+1$ *group (where* $0 \leq i < \mathcal{H}$*) contains a subgroup which is adjacent to all other subgroups of this group.*

These unique constraints and their corollaries [1] constitute invariants of the group hierarchy. Efficient maintainance of these invariants is possible even with such a simple traffic pattern as the periodic gossiping. Moreover, this approach allowed us to prove analytically that the protocol ensures consistency and convergence of the hierarchy after arbitrary changes in the node population or connectivity [1].

3 EVALUATION

We evaluated PL-Gossip using our own packet-level event-driven simulator. We conducted experiments with varying network sizes, densities, message loss rates, and node arrival and departure schemes [1]. The experimental results verified that a node's state, as maintained by the protocol (i.e., the label and the routing table), grows logarithmically with the network size, which ensures scalability and small bandwidth requirements. In addition, the hierarchical network organization offers efficient routing: the average hop stretch does not exceed 25%. Moreover, the hierarchy is quickly bootstrapped or restored, also under significant node population changes, which minimizes disruptions caused to the applications. Finally, the experiments confirmed what we proved analytically, that is, the protocol recovers the network from any massive node failure or network partitioning, even when such incidents happen continuously and concurrently.

We have also implemented PL-Gossip in TinyOS. The implementation is subject to real-world tests while at the same time being integrated into a large-scale real-world system.

References

[1] K. Iwanicki and M. van Steen, "The PL-Gossip algorithm," Vrije Universiteit, Amsterdam, the Netherlands, Tech. Rep. IR-CS-034, March 2007, available at: http://www.few.vu.nl/~iwanicki/.

[2] D. Ganesan, B. Greenstein, D. Estrin, J. Heidemann, and R. Govindan, "Multiresolution storage and search in sensor networks," *ACM Transactions on Storage (TOS)*, vol. 1, no. 3, pp. 277–315, August 2005.

[3] I. F. Akyildiz and I. H. Kasimoglu, "Wireless sensor and actor networks: research challenges," *Ad Hoc Networks (Elsevier)*, vol. 2, no. 4, pp. 351–367, October 2004.

[4] S. Kumar, C. Alaettinoglu, and D. Estrin, "SCalable Object-tracking through Unattended Techniques (SCOUT)," in *Proceedings of ICNP 2000*, Osaka, Japan, November 2000, pp. 253–262.

[5] K. Iwanicki and M. van Steen, "Towards a versatile problem diagnosis infrastructure for large wireless sensor networks," in *Proceedings of the 2nd OTM International Workshop on Pervasive Systems (PerSys 2007)*. Vilamoura, Portugal: Springer LNCS, November 2007, to appear.

[6] X. Li, Y. J. Kim, R. Govindan, and W. Hong, "Multi-dimensional range queries in sensor networks," in *Proceedings of SenSys 2003*, Los Angeles, CA, USA, November 2003, pp. 63–75.

[7] B. Greenstein, S. Ratnasamy, S. Shenker, R. Govindan, and D. Estrin, "DIFS: A distributed index for features in sensor networks," *Ad Hoc Networks (Elsevier)*, vol. 1, no. 2-3, pp. 333–349, 2003.

[8] D. Ganesan, D. Estrin, and J. Heidemann, "Dimensions: Why do we need a new data handling architecture for sensor networks?" *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 1, pp. 143–148, January 2003.

[9] J. Hagouel, "Issues in routing for large and dynamic networks," Ph.D. dissertation, Columbia University, 1983.