

VU Research Portal

Hierarchical Routing in Low-Power Wireless Networks

Iwanicki, K.S.

2010

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Iwanicki, K. S. (2010). *Hierarchical Routing in Low-Power Wireless Networks*. [PhD-Thesis - Research and graduation internal, Vrije Universiteit Amsterdam].

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Hierarchical Routing in Low-Power Wireless Networks

Wireless sensor networks (eng. *wireless sensor networks*, abbr. *sensornets* or *WSNs*) are novel, low-power wireless networks that constitute a new class of computing. They are composed of numerous tiny wireless devices that, embedded in physical space, can collaboratively sense and control various features of the surrounding environment. Their objective is thus to extend the digital world of the Internet with the capabilities of remotely sensing and controlling the physical world. To fulfill this task, such devices have to communicate wirelessly with each other and with other devices on the Internet.

The fundamental functionality that enables any two Internet devices to communicate is point-to-point routing. It is the only functionality implemented by all Internet devices, be they in the Internet core or at its edges. Since a device (a node) in the Internet can communicate directly only with a tiny subset of all nodes (e.g., a network cable connects just two nodes and a radio has only limited range), the data sent by the node has to be forwarded by some intermediate nodes in order to reach the destination node potentially on the other side of the globe; such intermediate nodes form a path. The objective of point-to-point routing is finding paths in a network along which data are sent between nodes. Therefore, point-to-point routing is crucial in the Internet to allow any two nodes to communicate, that is, to create an illusion of a full connectivity in a network in which few node pairs can communicate directly.

Although initially point-to-point routing was not deemed necessary for sensornets, recent research results have revisited that view: point-to-point routing is important for many sensornet applications, especially if sensor nodes are to be fully-fledged Internet devices. The goal of this dissertation is thus to develop a point-to-point routing protocol suitable for sensornets.

However, developing such a protocol is extremely challenging. Due to their embedding in physical space, sensor nodes are severely constrained in terms of resources, such as energy, memory, bandwidth, and processing power. Moreover, the low-power wireless communication they employ exhibits many peculiar phenomena that make it unreliable and hard to engineer. Finally, the nodes are typically deployed in large numbers to cover all sensing and control points.

Apart from some other minor challenges, the major challenge these features of sensornets introduce for routing is that any point-to-point routing protocol for sensornets has to simultaneously ensure the following properties. Firstly, considering the resource constraints of sensor nodes, a routing protocol has to offer small routing state, which is crucial for scalability and efficiency. Moreover, the protocol has to provide small routing stretch, that is, the routing paths it

finds should be close to the optimal ones, which is important for efficiency and reliability as it reduces the global resource consumption and improves end-to-end data delivery rates. Furthermore, the protocol has to be robust against communication and node failures, which is important for scalability and reliability, especially considering the unreliable nature of low-power wireless communication and the interactions of sensor nodes with the surrounding environment. Finally, the protocol should be self-managed to a large extent, as this simplifies deployment and maintenance of large networks, which are inherent in many sensornet applications.

These requirements are explained in more detail in Chapter 2, based on an analysis of existing and some envisioned sensornet applications. Moreover, the analysis constitutes a strong argument that they are unlikely to change with the technology progress in the near future. This is because, to be deployed more widely, sensornet hardware technology will continue emphasizing different aspects, such as low-power operation, form factor, and price per unit, leaving all the constraints and peculiarities of sensornets to be coped with in software. Consequently, there is every likelihood that the research results presented in this dissertation will remain relevant for extended periods of time.

An analysis of related work on point-to-point routing, as performed in the second part of Chapter 2, yields some crucial observations. On the one hand, there exists an entire spectrum of routing techniques that are potentially suitable for sensornets, promising robustness and self-management capabilities. On the other hand, theory on routing evidences that there is an inherent trade-off between routing state and routing stretch, which all such techniques have to make: in short, the smaller the state in a technique, the larger the stretch, and vice versa. Therefore, the spectrum of techniques can be divided into a few main regions depending on the particular state-stretch trade-off they involve. It can then be observed that one of the most promising regions in the techniques spectrum, the region corresponding to hierarchical routing, has not yet received much attention. In particular, unlike other representative techniques of the state-stretch trade-off spectrum, to the best of my knowledge, no hierarchical routing protocol has been implemented and evaluated in sensornets.

The principal idea behind hierarchical routing is to organize nodes into a multi-level hierarchy of clusters. With such an organization, instead of maintaining routing state for every other node in the network, a node can maintain state only for a few clusters in its vicinity. In this way, routing state can be reduced tremendously to a polylogarithmic function of the node population size.

However, such a reduction in routing state is traded off for an increase in routing stretch. In theory, the increase in stretch in hierarchical routing can be

large in some network topologies. It can be speculated that this is one of the main reasons for the scarcity of research on hierarchical routing for sensornets.

Another likely reason is that, even disregarding the difficulties stemming from the resource constraints and the peculiarities of sensornets, hierarchical routing is quite difficult to implement. Especially the fundamental problem of this routing technique — organizing nodes into the cluster hierarchy and maintaining the hierarchy when nodes fail or the network connectivity changes — is extremely complex. These two issues can make hierarchical routing potentially unappealing for sensornets, even despite its other merits.

This dissertation challenges this view and advocates hierarchical routing. There are three main motivating factors behind hierarchical routing. First, hierarchical routing offers very small routing state. Second, while the stretch in hierarchical routing can theoretically be large if arbitrary network topologies are considered, due to embedding in physical space, sensornet topologies are typically “geometric,” which means that a length of a routing path grows fast with growing node populations. Theory proved that in such topologies hierarchical routing can offer both small state and small stretch. Third, implementing hierarchical routing need not necessarily be more difficult than implementing other techniques. In particular, by relaxing some of the hierarchy properties, one can devise localized algorithms in which nodes self-organize into and autonomously maintain the hierarchy in a highly robust and efficient manner. Consequently, all in all, hierarchical routing has the potential to be a compelling point-to-point routing technique for sensornets.

To support the above argument, in Chapter 3, a practical hierarchical routing protocol for sensornets, called *PL-Gossip*, is introduced. To the best of my knowledge, *PL-Gossip* is the first hierarchical routing protocol that has been effectively implemented for sensornets. It decomposes hierarchical routing into three widely-recognized manageable abstractions corresponding to different aspects of routing in sensornets, namely link quality estimation, routing state maintenance, and packet forwarding, and explains how each of the abstractions can be implemented to support hierarchical routing.

The most challenging, albeit crucial routing abstraction in hierarchical routing is routing state maintenance, that is, the maintenance of the cluster hierarchy reflected in node routing tables and routing addresses. As mentioned above, the intricacy of this problem is likely one of the reasons that hierarchical routing has not yet been implemented and evaluated in sensornets. Therefore, *PL-Gossip* puts special emphasis on the cluster hierarchy maintenance problem.

To make the problem practically tractable, rather than aiming at optimal hierarchies, *PL-Gossip* focuses on best-effort ones. In particular, Chapter 3

illustrates how to customize the properties of a cluster hierarchy, such that it can be proved (cf. Appendix A) that the hierarchy: first, is suitable for hierarchical routing, second, has the potential to offer small routing state, and third, can be maintained by real sensor nodes. Such custom properties need not be always be the same as those in Chapter 3, though; in particular, they are changed in subsequent chapters of the dissertation, even to describe different hierarchy types. This is just one of the examples that, rather than being a fixed and monolithic protocol, *PL-Gossip* introduces ideas that are more broadly applicable.

To maintain the properties of a cluster hierarchy, the variant of *PL-Gossip* presented in Chapter 3 proposes to use a combination of two simple concepts: local operations for modifying the hierarchy and asynchronous local gossiping for propagating such modifications to the affected nodes. This is sufficient to allow nodes running *PL-Gossip* to self-organize into and autonomously maintain a cluster hierarchy described by a custom set of properties. However, simple though these two concepts are, combining them to enable self-managed hierarchy maintenance poses a number of challenges. Such challenges are associated, for example, with consistently adopting hierarchy modifications by the affected nodes or with electing those lower-level clusters that will be promoted to higher-level clusters. A lot of effort is dedicated to provide solutions to these non-trivial challenges and to prove the correctness of the solutions analytically. Again, however, these solutions are simple, which overall facilitates implementation of *PL-Gossip* on real sensor nodes. All in all, Chapter 3 illustrates that one can indeed devise a practical self-managed hierarchical routing protocol for sensornets.

The development of *PL-Gossip* enables evaluating hierarchical routing experimentally in various sensornet settings, which is the subject of Chapter 4. The evaluation employs three experimental platforms: a custom high-level simulator, which I have implemented, TOSSIM, a low-level simulator for sensor nodes with realistic models of low-power wireless communication, and an actual 50⁺-node testbed, which I have built at Vrije Universiteit (cf. Appendix B). Hierarchical routing is evaluated with respect to the aforementioned requirements for a sensornet routing protocol: routing state, routing stretch, and robustness.

The evaluation shows that, even though it does not aim at optimal cluster hierarchies, *PL-Gossip* can indeed offer very small routing state. Moreover, despite such small state, in practice *PL-Gossip* can also offer small stretch, on average within approximately 50% of the optimal one, which is the consequence of the aforementioned geometric nature of sensornets. Finally, *PL-Gossip* is also robust in that failures cause relatively little disruption and the protocol recovers from them relatively fast. These results suggest that hierarchical routing is indeed

an appealing point-to-point routing technique for sensor networks.

Furthermore, the evaluation of hierarchical routing in Chapter 4 is not only limited to the presented *PL-Gossip* protocol. As argued above, rather than being a fixed and monolithic protocol, *PL-Gossip* constitutes a basis of a protocol and introduces ideas that are more broadly applicable and can be experimented with. Therefore, rather than only implementing *PL-Gossip*, an entire hierarchical routing library has been developed for the evaluation. The library defines a common framework for a hierarchical routing protocol for sensor networks, which draws from the lessons learned from *PL-Gossip*. At the same time, however, it identifies a number of design points. By varying the solutions at these design points, one can evaluate and compare different mechanisms proposed for hierarchical routing, as well as some novel ideas. The source code of the library has been made publicly available (cf. Appendix C.1).

The library enables conducting thorough evaluation of different design decisions that affect routing state, routing stretch, and robustness of hierarchical routing itself. One of such design points studied in Chapter 4 is the type and the properties of a cluster hierarchy. In particular, it is shown that a landmark hierarchy typically requires larger state than an area hierarchy, but in return offers smaller stretch, on average within approximately 25% of the optimal one; this can be further controlled by varying particular properties of the hierarchy. In other words, by varying the type or the properties of a cluster hierarchy, one can further explore the state-stretch trade-off within hierarchical routing itself. Another design point studied in Chapter 4 is the method for propagating hierarchy information. Apart from local asynchronous gossiping, adopted by the variant of *PL-Gossip* from Chapter 3, a common flooding-based method and a novel hybrid one are studied. In short, it is shown that different methods affect different aspects of robustness, thus again, hierarchical routing can be customized to optimize a particular metric of interest. All in all, not only does the evaluation of hierarchical routing conducted in Chapter 4 confirm that hierarchical routing is appealing for sensor networks, but also shows that it can be customized for particular applications.

Since the research presented in Chapter 3 and Chapter 4 fills in the aforementioned gap in the routing techniques spectrum, it makes it possible to experimentally compare representative techniques from the entire spectrum. Such an initial comparison is the subject of Chapter 5.

To perform the comparison, another routing library, this time covering various techniques, has been implemented, as described in Chapter 5. The library involves four such techniques, which together represent the entire state-stretch trade-off spectrum: shortest-path routing, which requires the largest state but offers the minimal stretch, compact and hierarchical routing, which at a different granularity

reduce state at the expense of stretch, and constant-state routing, which needs the smallest state but delivers the largest stretch. The great majority of the code is shared by all techniques, which makes the implementations of these techniques uniform to a large extent. The library is evaluated in TOSSIM and on two testbeds: a 50⁺-node testbed and a 100⁺-node testbed. The evaluation focuses on the state-stretch trade-off.

Apart from some minor results, the evaluation makes two major contributions. First, it illustrates a general property of sublinear-state routing techniques. More specifically, due to the geometric nature of sensornet deployments, most techniques with sublinear state offer a routing stretch that is close to the optimal one. In other words, in sensornets, a large reduction in routing state typically entails only a slight increase in routing stretch. This leads to the second contribution, which is directly related to hierarchical routing. More specifically, compared to the other representative techniques, hierarchical routing offers a large reduction in state with only a slight increase in stretch. In particular, compared to compact routing, hierarchical routing offers a state that allows for at least an order of magnitude better scalability, while compromising the stretch only slightly, by less than 10–15% on average, which for many sensornet applications will degrade performance insignificantly. All in all, apart from reinforcing the argument that hierarchical routing is an appealing technique for sensornets, Chapter 5 suggests that, from the representative techniques from state-stretch trade-off spectrum, hierarchical routing offers a trade-off that is arguably most appealing for many sensornet applications.

To sum up, the overall conclusion that can be drawn from all the research results is that the thesis formulated in this dissertation holds.

Hierarchical routing is a compelling point-to-point routing technique for large sensornets. In practice, not only does it offer small routing state, but also small routing stretch. Moreover, it is possible to provide robust, efficient, self-managed hierarchical routing protocols that work in the real world.