

Delivery Semantics for Geographic Routing

Matthias Witt and Volker Turau

Hamburg University of Technology, Institute of Telematics
Schwarzenbergstraße 95, 21073 Hamburg, Germany

Abstract. Geographic routing algorithms use locations rather than node addresses as destinations. But since it is not always guaranteed that there is indeed a node exactly at the destination location, there must exist a description which nodes shall be allowed to consume a message. In some cases, the message must not be consumed by a node at another location (e. g., if the sink is the destination), whereas in other cases, nodes in the vicinity may also consume it. This has to be chosen by the application, so the routing protocol should support different delivery semantics. Another question is if only one node may consume the message, or if multiple destination nodes are allowed.

This paper proposes three orthogonal delivery parameters: closeness, multiplicity, and accept-outside.

1 Introduction

Geographic routing is an important technique for wireless sensor networks, since it avoids the storage of routing tables or path information in the sensor nodes, which are usually highly resource-constrained. The presence of information about the node location (either through GPS or some localization algorithm) and—for some algorithms—the direct neighbors is sufficient.

Several geographic routing algorithms for sensor networks have been proposed. They use different, but elaborated techniques to forward packets using location information. One important question, however, that is often neglected is to describe the conditions under which a node is allowed to declare itself as destination and consume the message. Since the destination is not a node address but a location, this is substantially different from traditional routing schemes. Many algorithms assume that there exists a node exactly at the destination location [1, 2], the node locations being published for instance by a distributed location service like GLS [3]. The location service, however, imposes an extra overhead. Besides, the destination locations may be incorrect due to location errors [4]. Additionally, nodes may fail or move, so it is not always possible to keep track of the current state. Other algorithms perform a limited flooding in the vicinity of the destination [5], which also requires the sending of many packets. An interesting alternative is utilized by GHT [6], which takes advantage of GPSR's perimeter mode: The destination location is traversed using the right-hand rule; when the message has looped around the location, the first node that receives the message for the second time in perimeter mode declares itself as destination and consumes the message.

Although there are different strategies for finding a suitable destination node, each algorithm uses only a single strategy. It is, however, dependent on the specific application if, for example, the destination's exact location is known by the sender or if the message should be delivered "somewhere" in the vicinity of a specific location. In some cases, it is important that *at most* one node consumes the message, whereas in other cases it is more important that *at least* one destination node is found at all.

Anycast is a well-known technique for routing messages to any out of several possible destinations. Within the scope of sensor networks, anycast can be used when multiple sinks exist; messages are routed to the nearest sink [7–9]. However, anycast routing schemes do not use geographic routing and consequently are highly application-specific.

Other algorithms include time in delivery semantics. *Mobicast* [10] is a spatiotemporal multicast scheme that supports moving destination areas. Since the semantics proposed in this paper are not dependent on time, such schemes are not covered.

This paper provides a taxonomy of delivery semantics and proposes three parameters, which can be combined independently. An application that sends messages can set these parameters according to the specific requirements. A concrete routing algorithm must then deliver the message with respect to the desired semantics. The proposed delivery semantics are independent from concrete routing algorithms. The implementation of the semantics is beyond the scope of this paper.

2 Delivery Semantics

Basically, there are three questions that have to be answered when describing the geographical destination for a message in a sensor network:

1. How close to the destination must a sensor node be in order to consume the message?
2. Is it acceptable that multiple nodes consume the message, or should only one node consume it?
3. When the message gets stuck because there is no node that is close enough to the destination, should the node where the message got stuck drop the message or consume it if the node is sufficiently close to the destination?

In the following, these questions are discussed in detail.

2.1 Closeness

Assume a sensor network for environmental monitoring, where neither the user nor the base station know the exact positions of the nodes. The user requests the sensor data for a specific location and initiates a query, having this location as destination position. However, he does not know whether there is a sensor node exactly at this location. Moreover, this is extremely improbable. Hence, it

should be acceptable that a node consumes the message if its distance to the destination location lies within a designated limit of tolerance t . Using Euclidean metrics, the destination area is a circle with radius t . This approach raises the probability that a destination node is found. The limit can be set depending on the network density. Assume that the network consists of n nodes spread uniformly in a field of surface area A . Then the probability that the destination area is empty is

$$\left(1 - \frac{\pi t^2}{A}\right)^n,$$

not regarding border effects. If this probability should be less than p , t has to be set to a value such that

$$t > \sqrt{\frac{A}{\pi} \left(1 - p^{\frac{1}{n}}\right)}.$$

Consider, on the other hand, the case that the measurement message is sent from the sensor node to the sink that initiated the query. The message is explicitly addressed to the sink, so it is known that there is a node at the designated location. In this case, the limit of tolerance can be set to zero; this also ensures that no other node than the sink consumes the message.

These examples show that at least two different semantics are needed for describing destination locations: *exact* and *nearby*. In the latter case, a limit of tolerance has to be given, in the former one, it is zero.

There may be cases in which a circle is not appropriate as destination area, e. g., when the nodes are distributed in a regular mesh.

2.2 Multiplicity

If the *nearby* semantics is used, there may be more than one potential destination node. The routing protocol may take care that only one of these nodes consumes the message; however, this leads to more overhead, since the nodes must somehow agree about the winner, which comes at additional communication costs. But this is not necessary in cases where it is not important that only one node receives the message. For example, if the sensor node at (or nearby) a specific location shall be set in an alarming mode, because the user wants to have special attention to this location, it is not crucial that only one node receives the message. To account for node failures, it is even better if multiple nodes receive it.

On the other hand, there are cases where only one node shall receive the message. In general, this is the case when the nodes serve different tasks. For instance, when the message is addressed to the sink, no other node is allowed to consume it.

Inspired by RPC semantics, the following semantics are possible for the multiplicity of the destination:

maybe: the message may reach a node or not,

exactly-one: the message must be consumed by exactly one node,

at-most-one: the message must be consumed by zero or one node,
at-least-one: the message must be consumed by one or more nodes,
all: the message must be consumed by all nodes meeting the closeness semantics.

Note that these semantics are not meant for describing *how many times* a node receives the message. This aspect of delivery semantics is orthogonal to the ones discussed here. The semantics specify *how many nodes* receive the message.

2.3 Accept-outside

Consider the case that there is no node that matches the destination description, but the message has reached a node that is already near the destination. The destination location is within the transmission range, the current node does not yet match the destination description, but no closer node can be found, either because closer nodes do not exist or because they are currently not reachable. Now, there are two possibilities: Either the message is regarded as non-deliverable and dropped, or the node where the message got stuck declares itself as destination and consumes the message. The strategy to be applied depends on what has higher priority: that a node which exactly matches the destination description receives the message, or that actually any node receives it at all. For example, if data at a specific location shall be measured, it may be tolerable if merely data near this location are measured; however, if the sink is the destination, it is not acceptable when another node consumes the message. This demands for a Boolean parameter *accept-outside*.

This parameter can be used in combination with *exact* semantics to deliver the message to the node that is closest to a location. This is not possible with *nearby* semantics, since in this mode the message is consumed by any node within the area of tolerance, not necessarily by the closest one.

Not all combinations of these parameters seem to be useful. For instance, in *exact* semantics with *accept-outside=false*, the semantics *at-least-one* and *exactly-one* seem to yield the same results (when there are not several nodes with identical positions). However, the routing protocol may act differently; for example, it may deliver the messages faster or send fewer packets in *at-least-one* semantics. Admittedly, the concrete implementation of the semantics is beyond the scope of this paper.

An important issue is that for some semantics it is not guaranteed that the same node(s) consume(s) subsequent messages sent to the same location. The messages may be forwarded on different paths and hence reach different nodes. Only using *exact* semantics or *all* semantics, with *accept-outside=false* in both cases, it can be guaranteed that all messages are consumed by the same node(s).

3 Conclusion

For geographic routing algorithms it is essential to include a clear description of the destination. Different applications have different requirements regarding the node or nodes that shall consume the messages. Therefore, the routing protocol must support different delivery semantics. In this paper, three orthogonal parameters have been proposed: *closeness*, *multiplicity*, and *accept-outside*.

Implementations of these semantics are currently investigated by extending selected geographic routing algorithms.

References

1. Karp, B., Kung, H.T.: GPSR: Greedy Perimeter Stateless Routing for Wireless Networks. In: Proceedings of the 6th Annual International Conference on Mobile Computing and Networking, Boston, Massachusetts (2000) 243–254
2. Füßler, H., Widmer, J., Mauve, M., Hartenstein, H.: A Novel Forwarding Paradigm for Position-Based Routing (with Implicit Addressing). In: IEEE Computer Comm. Workshop (CCW 2003), Dana Point, California (2003) 194–200
3. Li, J., Jannotti, J., De Couto, D., Karger, D., Morris, R.: A Scalable Location Service for Geographic Ad Hoc Routing. In: Proceedings of the 6th Annual International Conference on Mobile Computing and Networking, Boston, Massachusetts (2000) 120–130
4. Witt, M., Turau, V.: The Impact of Location Errors on Geographic Routing in Sensor Networks. In: Proc. International Conference on Wireless and Mobile Communications (ICWMC'06), Bucharest, Romania (2006)
5. Heissenbüttel, M., Braun, T., Bernoulli, T., Wälchli, M.: BLR: Beacon-Less Routing Algorithm for Mobile Ad-Hoc Networks. Elsevier's Computer Comm. Journal **27**(11) (2004) 1076–1086
6. Ratnasamy, S., Karp, B., Shenker, S., Estrin, D., Govindan, R., Yin, L., Yu, F.: Data-Centric Storage in Sensornets with GHT, A Geographic Hash Table. Mobile Networks and Applications **8**(4) (2003) 427–442
7. Intanagonwiwat, C., De Lucia, D.: The Sink-based Anycast Routing Protocol for Ad Hoc Wireless Sensor Networks. Technical Report 99-698, Computer Science Department, University of Southern California (1999)
8. Thepvilajanapong, N., Tobe, Y., Sezaki, K.: HAR: Hierarchy-Based Anycast Routing Protocol for Wireless Sensor Networks. In: Proc. 2005 IEEE/IPSJ Int. Symposium on Applications and the Internet (SAINT'05), Trento, Italy (2005) 204–212
9. Hou, Y.T., Shi, Y., Sherali, H.D.: Optimal Base Station Selection for Anycast Routing in Wireless Sensor Networks. IEEE Trans. on Vehicular Technology **55**(3) (2006)
10. Huang, Q., Lu, C., Roman, G.C.: Spatiotemporal Multicast in Sensor Networks. In: Proc. 1st Int. Conference on Embedded Networked Sensor Systems (SenSys'03), Los Angeles, California (2003) 205–217