

Hierarchical Geographic Multicast Routing for Wireless Sensor Networks

Dimitrios Koutsonikolas¹, Saumitra Das¹, Y. Charlie Hu¹, and Ivan Stojmenovic^{2,3}

¹*School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907*

²*Electronic, Electrical & Computer Engineering, The University of Birmingham, United Kingdom*

³*SITE, University of Ottawa, Ontario, Canada*

Abstract

Multicast is a fundamental routing service for efficient data dissemination required for activities such as code updates, task assignment and targeted queries in large-scale wireless sensor networks. Recently, two protocols were proposed to optimize two orthogonal aspects of location-based multicast protocols: GMR [1] improves the forwarding efficiency by exploiting the wireless multicast advantage but it suffers from scalability issues when dealing with large sensor networks; HRPM [2] reduces the encoding overhead by constructing a hierarchy at virtually no maintenance cost via the use of geographic hashing but it is energy-inefficient due to inefficiencies in forwarding data packets. In this paper, we present HGMR (Hierarchical Geographic Multicast Routing), a new location-based multicast protocol that seamlessly incorporates the key design concepts of GMR and HRPM and optimizes them for wireless sensor networks by providing both forwarding efficiency (energy efficiency) as well as scalability to large networks. Our simulation studies in an ideal environment and in a realistic environment confirm that HGMR exhibits the strength of both GMR and HRPM.

1. Introduction

In wireless sensor networks, *multicast* is a fundamental routing service for efficient data dissemination required for activities such as code updates, task assignment and targeted queries. In particular, *efficient* multicast for sensor networks is particularly critical due to the limited energy availability in such networks. Multicast protocols that exploit location information available from GPS [3] or localization algorithms are more efficient and robust than other stateful protocols. Since localization is typically already required for sensing applications, this location information can simply be reused for optimizing multicast performance at no extra cost.

Several location-based multicast protocols for wireless networks (not specifically sensor networks) have been proposed [4, 5, 6] which neither assume any unicast routing

scheme nor build any distributed multicast routing structure. These protocols build multicast trees using location information and use geographic forwarding to forward packets down the multicast trees. These protocols are *stateless*, as they carry encoded membership and location as well as tree information in each packet, so that the multicast membership and routing state do not have to be distributed as in traditional multicast protocols such as MAODV [7], ADMR [8] and ODMRP [9]. Stateless protocols are more robust and potentially more efficient than stateful protocols as they avoid the difficulty of maintaining distributed state in the presence of frequent topology changes due to node failure or mobility.

Despite the high promise of the location-based protocols, there are two major challenges to further improvement of these class of protocols. First, these stateless protocols typically use geographic forwarding for data dissemination down each branch of the multicast tree, and hence they are not exploiting the multicast advantage of the wireless transmission. Second, as the size of the network increases, even for a sparse membership density for which these stateless protocols are particularly good for, the overhead of the encoding membership in each data packet will become significant. Most recently, two protocols were independently proposed to address the above two orthogonal challenges respectively. GMR [1] improves the forwarding efficiency of location-based multicast by exploiting the wireless multicast advantage (WMA), while HRPM [2] reduces the encoding overhead of location-based multicast by constructing a hierarchy at virtually no maintenance cost via the use of geographic hashing.

In this paper, we present HGMR (Hierarchical Geographic Multicast Routing), a new location-based multicast protocol that seamlessly incorporates innovations in location-based multicast and optimizes them for wireless sensor networks by providing both encoding efficiency (and hence energy-efficiency) as well as scalability to large networks. HGMR starts with a hierarchical decomposition of a multicast group into subgroups of manageable size (i.e. encoding overhead) using HRPM's key concept *mobile geographic hashing*. Within each subgroup, HGMR

uses GMR’s *local multicast scheme* to forward a data packet along multiple branches of the multicast tree in one transmission. Thus, HGMR can simultaneously achieve energy-efficiency (through higher forwarding efficiency utilizing multicast advantage) and scalability (through low overhead hierarchical decomposition).

We first evaluate the performance of HGMR by comparing it to GMR and HRPM in an ideal environment, in order to isolate the impact of the environment from the raw performance of different protocols and hence get a clear picture of the forwarding efficiency and the encoding overhead of each protocol. Our evaluation results show that with a large network size of 1000 nodes, HGMR incurs fewer transmissions (13% less) even when compared to GMR, while maintaining a low encoding overhead of 15%, very close to that of HRPM (12%) and much lower than that of GMR (36%).

We then compare the performance of HGMR with those of GMR and HRPM in a realistic environment, in order to take into account the effect of practical factors such as loss and contention and to evaluate the performance of the protocols from the real applications’ perspective (by looking at metrics such as PDR and packet delivery latency). Our evaluation results show that HGMR outperforms the other two protocols in terms of both of these metrics. As the network size increases to 1000 nodes, HGMR maintains a PDR above 85% and a latency lower than 52 msec, similarly as HRPM though HRPM has 83% more transmissions, while GMR’s performance deteriorates, resulting in a low PDR of only 60% and a high packet delivery latency (23% higher than HRPM’s and 27% higher than HGMR’s).

2. Hierarchical Geographic Multicast Routing

In this section, we first give a brief overview of GMR [1] and HRPM [2]. We then present the design of HGMR.

2.1. Overview of GMR and HRPM

GMR Geographic Multicast Routing (GMR) [1] exploits the wireless multicast advantage to improve the forwarding efficiency of previous location-based (and hence stateless) multicast protocols. Similarly as in previous protocols, it assumes centralized membership management at the multicast root. Differently from in previous protocols, each node along the multicast tree tries to send a data packet down multiple branches of the multicast tree using one broadcast transmission.

In more detail, each forwarding node propagating multicast data needs to select a subset of its neighbors as relay nodes towards destinations. With GMR, nodes make this selection by optimizing the cost over progress ratio. The *cost* is equal to the number of selected neighbors, while the *progress* is the overall reduction of the remaining distances to the destinations. Such a neighbor selection scheme

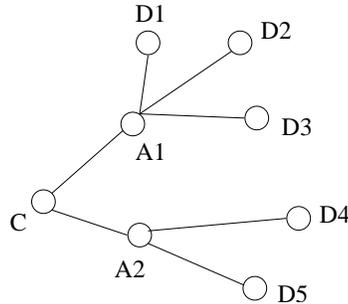


Figure 1. Forwarding node selection in GMR

achieves a good tradeoff between the cost of the multicast tree and the effectiveness of data distribution.

We explain the cost over progress ratio metric with the help of Figure 1. In this figure, assume that node C, after receiving a multicast message, is responsible for destinations D_1, D_2, D_3, D_4, D_5 , and it is considering its neighbors A_1 and A_2 as possible relay nodes. The current total distance for multicasting is $T_1 = |CD_1| + |CD_2| + |CD_3| + |CD_4| + |CD_5|$. If C selects A_1 as the relay node responsible for D_1, D_2, D_3 and A_2 as relay node responsible for D_4 and D_5 , the new total distance is $T_2 = |A_1D_1| + |A_1D_2| + |A_1D_3| + |A_2D_4| + |A_2D_5|$ and the progress made is $T_1 - T_2$. Hence, the cost over progress ratio for the forwarding set $\{A_1, A_2\}$ is $\frac{2}{T_1 - T_2}$.

In the implementation, GMR adds a header to data messages to allow the selected neighbors to realize they are selected as relay nodes. Each node that forwards a message includes in the header the IDs of the neighbors it has selected as relay nodes and the IDs of the destinations each of the selected relay nodes is responsible for. The message is then *broadcast* and it can be received by all of its neighbors due to the broadcast nature of the wireless channel. This reduces the total number of transmissions and hence the total energy consumed, since each node along the multicast tree only needs to send a single message in order to deliver a multicast data packet to multiple relay nodes.

HRPM Hierarchical Rendezvous Point Multicast (HRPM) [2] reduces the encoding overhead by employing two key design concepts: (1) Use of hierarchical decomposition of multicast groups and (2) Leveraging geographic hashing to construct and maintain such a hierarchy efficiently. Figure 2 shows the major components of HRPM.

The main design goal of HRPM is to limit the per-packet overhead to an application-specified constant (ω), irrespective of the group size G. It achieves this by recursively partitioning a large multicast group into manageable-sized subgroups in which the tree-encoding overhead satisfies the ω constraint. This partitioning is achieved by geographically dividing the deployment area into smaller and smaller cells,

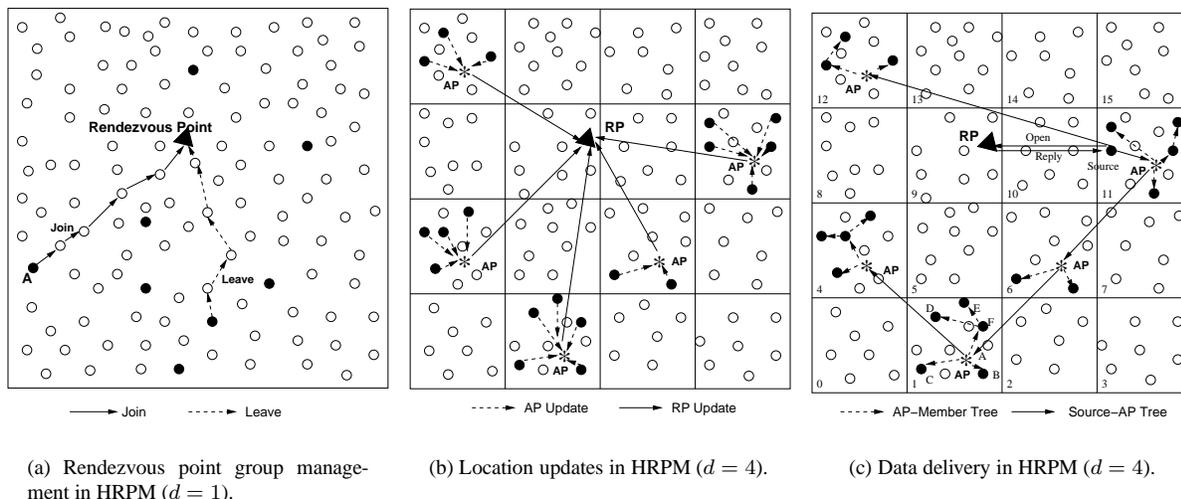


Figure 2. Group management, tree construction, and data delivery in HRPM.

which form a hierarchy with the root representing the entire region. Every cell in the hierarchy has an AP (Access Point) and the entire region has an RP (Rendezvous Point). All members in a leaf cell of the hierarchy form a subgroup and are managed by that cell's AP. Groups of APs are managed recursively, i.e., by the APs of their parent cells. Finally, APs belonging to the the highest level of the hierarchy are managed by the RP. [2] showed that a 2-level hierarchy is enough even for very large-size groups (up to 5800 group members). In this case, the area is divided in d^2 cells, each with one AP, and the d^2 APs are managed by the RP. The parameter d is called *decomposition index*.

To avoid the need of keeping track of the AP/RP nodes, needed in both membership management and data dissemination, which would require an external location service, HRPM adopts the idea of *geographic hashing* to reduce the maintenance of AP/RP nodes at virtually no maintenance cost. The role of each AP and as well as the RP is mapped to a unique geographic location, vis some simple hash function, the node that is currently closest to that location then serves the role of AP/RP, and routing to the AP/RP is conveniently achieved via geographic forwarding. Finally, similarly as geographic routing to a node, routing to a hashed location can also encounter *holes* in the network topology. HRPM [2] solves the problem using a technique similar to *face-routing* [10, 11, 12].

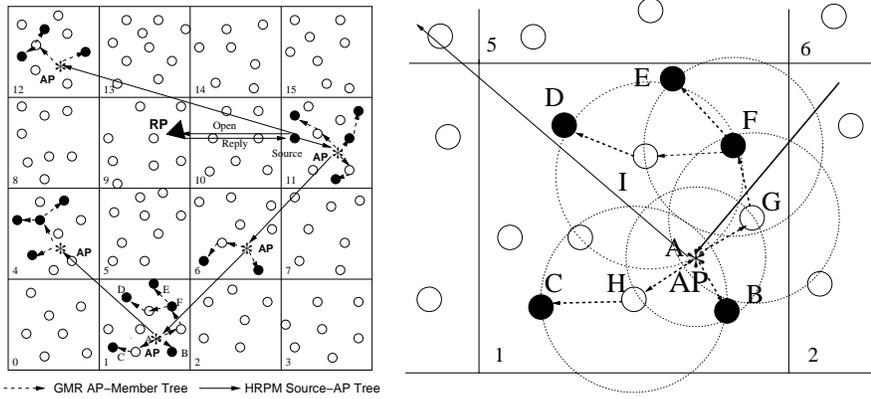
When a source has data packets to send, it first hashes the multicast group's identifier to obtain the location of RP. It then contacts the RP and obtains the group membership vector, which specifies which cells have members (or in other words which of the APs are active). After that, the source builds an overlay tree, the *Source* \rightarrow *APs* tree, considering each active AP as a vertex in a topology graph, and it sends data packets along the branches of this tree, using geographic forwarding. Each AP also builds an overlay tree,

the *AP* \rightarrow *Members* tree, whose vertexes are the members in that AP's cell and forwards the data along this tree, also using geographic forwarding. Both the source and the APs use unicast to forward data packets, which means the same packet is sent over each branch of a subtree, with a different header, specifying the paths it should follow along that subtree.

2.2. Hierarchical Geographic Multicast Routing (HGMR)

The overview above shows the strength and weakness of GMR and HRPM protocols. On one hand, GMR reduces the number of transmissions required to send a multicast data packet from the source to all destinations, since at each hop the packet is broadcast to all neighbors. However, this means that information about all the destinations and all the selected relay neighbors has to be included in the packet header. Assuming a reasonable group member density, as the size of the network increases, the byte overhead associated with each packet may increase to unacceptable levels. On the other hand, HRPM efficiently reduces the byte overhead associated with each data packet by dividing a large group into multiple subgroups. However, HRPM is inefficient in terms of packet transmissions, since at each node along the *Source* \rightarrow *APs* or the *AP* \rightarrow *Members* tree, the same data packet is unicast to possibly more than one neighbor nodes (each connecting to one subtree). Unicasting the same data packet more than once not only consumes bandwidth, which is limited in wireless networks, but also exhausts faster the nodes energy which is limited, especially in sensor networks.

In this section, we propose Hierarchical Geographic Multicast Protocol (HGMR) which seamlessly combines the scalability (low encoding overhead) of HRPM with the



(a) Data Delivery in HGMR ($d = 4$). Both the HRPM *Source* \rightarrow *APs* overlay tree and the GMR *AP* \rightarrow *Members* trees are shown. The *AP* \rightarrow *Members* trees connect both members (black dots) and non-members (white dots), as opposed to the overlay trees in Figure 2(c).

(b) *AP* \rightarrow *Members* tree construction in cell 1 using GMR's neighbor selection algorithm.

Figure 3. Data delivery in HGMR. Overall picture (a) and localized neighbor selection scheme (b).

forwarding efficiency of GMR.

HGMR divides the multicast group into subgroups using the mobile geographic hashing idea proposed in HRP: the deployment area is again divided into a number of cells, and in each cell there is an AP responsible for all members in that cell. When a source has data packets to send, it uses HRP's forwarding strategy to send the packets to each AP along the *Source* \rightarrow *APs* overlay tree. But within each cell, instead of constructing a *AP* \rightarrow *Members* overlay tree, HGMR uses GMR's cost over progress optimizing broadcast algorithm to select the next relay nodes at each hop. Adjusting the value for the decomposition index d , we can always ensure that the number of members an AP is responsible for does not grow too large. Hence, the use of GMR within each cell instead of HRP's unicast-based forwarding strategy helps to reduce the number of transmissions while maintaining a low encoding overhead.

The operations of HGMR are shown in Figures 3(a), 3(b). Figure 3(a) shows the overall picture, which includes the geographic division of the deployment area into cells, the *Source* \rightarrow *APs* overlay tree (same as the one constructed by HRP in Figure 2(c)), and the *AP* \rightarrow *Members* trees, one within each cell that contains some destinations. These trees are constructed using GMR's localized neighbor selection algorithm, and hence they are not overlay trees, as opposed to in HRP. Figure 3(b) shows in more detail the *AP* \rightarrow *Members* tree construction in cell 1. Here, the AP node *A* initially selects node *H* as relay node responsible for destination *C*, and node *G* as relay node responsible for destinations

D, *E*, *F*, and it can also reach directly destination *B*. Hence it sends only one message, with all this information to all its three neighbors, *H*, *G*, and *B*, exploiting the WMA. Node *B* finds that it is not responsible for any other destination, hence it does not rebroadcast the message. Node *H* rebroadcasts it to destination *C*. Node *G* can reach directly destination *F* and it also selects it as a relay for destinations *E* and *D*. *F* can reach *E* directly and it selects node *I* as a relay for destination *D*. Finally, *I* sends the message to *D*. There are a total of 5 data transmissions although the multicast tree has 8 edges.

3. Evaluation in an Ideal Environment

In this section, we evaluate the performance of GMR, HRP and HGMR in an ideal environment, i.e., with no packet loss.

3.1. Methodology

We implemented GMR, HRP, and HGMR in Glosim [13], a widely used wireless network simulator with a detailed and accurate radio model. To simulate ideal conditions, we used the free space propagation model and modified the MAC layer to remove collisions. Under this setup, all packets multicast by the source are properly received by all the destinations. We implemented HRP following [2] and GMR following [1]. In all the sections HRP adjusts the decomposition index d to the group size based on the equations in [2], using $\omega = 20\% \times PacketSize$. The same d is then used for HGMR.

We evaluated the scaling properties of the three protocols as a function of the network size. We varied the number of sensor nodes between 100 and 1000, but we kept the same node density (by varying the area of the deployment field), 20 nodes per radio range, and the same member density (30% of the nodes). Hence, as the network size increases, the group size increases as well. In all cases, the simulation duration was 600 sec. We had one multicast source that sent 512-byte packets at a constant rate of 2 packets/sec.

Evaluation metrics The following metrics are used to evaluate the efficiency of the three multicast protocols: (i)*Number of transmissions*: The total number of transmissions used to deliver the packets from the source to all the destinations. It measures the efficiency of the multicast paths selected. (ii)*Percentage of Forwarding Nodes (FNs)*: The number of nodes (including source) that transmitted at least one data packet divided by the total number of nodes. (iii)*Normalized Encoding Overhead (NEO)*: The ratio of the total number of encoding bytes transmitted at every hop to the total number of data bytes transmitted at every hop. Encoding bytes are the bytes used in each data packet to encode the information required by each protocol.

3.2. Results

Figures 4(a), 4(b), and 4(c) show the number of transmissions, the percentage of FNs, and the NEO for the three protocols with different network sizes. We make the following observations.

Figure 4(a) shows that the number of transmissions increases with the network size for all three protocols, but HRPM has the steepest increase. In the case of 1000 nodes, HRPM incurs 59% more transmissions than GMR and 72% more transmissions than HGMR. GMR uses MAC layer broadcast to transmit packets, exploiting WMA, while HRPM uses MAC layer unicast. Hence at each node fork of the multicast tree, only one transmission is required for GMR, and these transmissions can be heard by all 1-hop neighbors. On the other hand, with HRPM a node has to transmit the same packet separately to each of its neighbors selected as forwarding nodes. HGMR uses MAC layer unicast along the *Source* \rightarrow *APs* tree and MAC layer broadcast along the *AP* \rightarrow *Members* trees. Since there are fewer transmissions in the higher level tree compared to in the lower level trees, even though the higher level tree uses unicast, HGMR's total number of transmissions is still much lower than HRPM's and only slightly higher than GMR's for most network sizes.

Interestingly, for very large (1000 nodes) network sizes GMR is no longer the most efficient protocol in terms of the number of transmissions, incurring 8% more transmissions than HGMR. The large network size increases the lengths of the paths from the source to the destinations and GMR's FN selection algorithm selects more nodes, in trying to in-

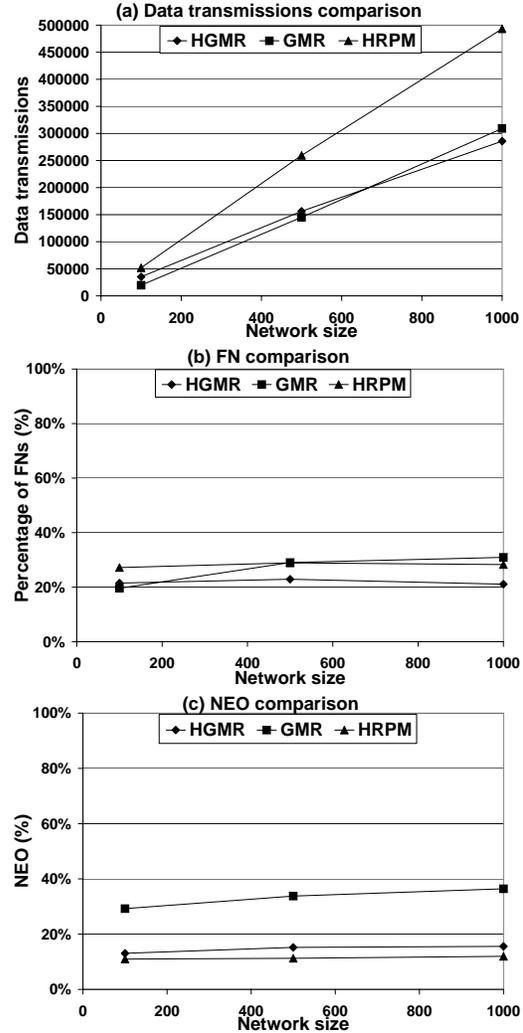


Figure 4. Performance comparison of GMR, HGMR and HRPM for different network sizes in an ideal environment.

crease the progress ratio, thus deteriorating forwarding efficiency. On the other hand, HGMR tries to balance the size of the *Source* \rightarrow *APs* and *AP* \rightarrow *Members* trees, and hence for *AP* \rightarrow *Members* trees, GMR's algorithm is never applied to a very large tree. This reduces the number of transmissions for HGMR compared to GMR.

Figure 4(b) shows that, as the network size increases, the percentage of FNs with HRPM becomes higher than HGMR's by about 7%. Another interesting observation here is that the percentage of FNs remains constant for the two hierarchical protocols, HRPM and HGMR, but it increases with the network size for GMR. Hence, although with a network of 100 nodes, GMR has the lowest percentage of FNs (19% vs 21% for HGMR and 27% for HRPM), with 1000 nodes, it has the highest one (31%). This is some-

what surprising, since GMR’s FN selection algorithm explicitly tries to minimize the number of FNs at each (real) routing hop, as opposed to HRPM which constructs minimum spanning overlay trees to deliver packets to the destinations (the *Source* \rightarrow *APs* and *AP* \rightarrow *Members* trees) and uses geographic routing to route packets between any two nodes of an overlay tree. However, it agrees with the result for the number of transmissions we analyzed before, and it shows again that GMR’s FN selection algorithm is not very efficient with very large network sizes.

Finally, from Figure 4(c) we observe that NEO increases with the network size for GMR, while for HRPM and HGMR it remains unaffected and much lower compared to GMR. With 1000 nodes, NEO is 12% for HRPM and 15% for HGMR but it reaches 36% for GMR. By increasing the network size, the number of multicast members also increases with a constant member density. With GMR a node sends the same data packet to all its neighbors, and each packet includes a list of all the neighbors selected as forwarding nodes as well as the information about all the destinations the selected neighbors are responsible for. As the group size increases, the source and nodes near the source have to include information about more and more destinations in the data packets they send – in the case of 1000 nodes with 300 members, the source has to include all the 300 members in the packets it sends. This results in a very large NEO for GMR. Although this increase in NEO has no effect in an ideal environment, it has a severe impact on GMR’s performance in a realistic environment, as we will see in Section 4. HRPM and HGMR on the other hand can adjust the number of members within each subgroup by varying the parameter d and hence, they maintain a low encoding overhead for all network sizes.

In summary, we saw that HRPM incurs too many packet transmissions, when the network size increases, and it also involves many nodes in packet forwarding. Hence, HRPM is not appropriate for sensor networks, where energy conservation is of great importance. In contrast, GMR incurs much fewer packet transmissions compared to HRPM, and it also uses fewer FNs, by exploiting WMA. However, for very large network sizes, the greedy neighbor selection of GMR is not as efficient and the percentage of FNs increases, increasing the energy consumption and limiting network’s lifetime. In addition, GMR has much higher encoding overhead compared to HRPM, and this will affect its performance in realistic scenarios. HGMR combines the high forwarding efficiency of GMR with the low encoding overhead of HRPM, and hence scales well as the group size increases with the network size.

4. Evaluation in a Realistic Environment

In this section we evaluate the performance of GMR, HRPM and HGMR in a realistic environment.

4.1. Methodology

To simulate a realistic environment, in Glomosim we used an IEEE 802.11 radio with a bit rate of 2Mbps and a transmission range of 250m. The *TwoRay* propagation model was used instead of the *Free Space*. Under this setup, packet loss can happen due to two reasons: (i) Collisions may happen when two or more nodes transmit a packet at the same time. The probability of collisions increases with the number of forwarding nodes and the packet size. (ii) Packets may be corrupted due to noise or the receiver may be unable to decode them due to low SNR. The probability of packet corruption increases with the packet size. All other simulation parameters remain the same as in Section 3.

Evaluation metrics In addition to the metrics used in Section 3, we used the following evaluation metrics: (i)*Average Packet Delivery Ratio (PDR)*: The number of the data packets delivered to a multicast group member divided by the number of data packets transmitted by the source, averaged over all multicast group members. This metric is necessary, since in a realistic environment, there is packet loss, as we mentioned above. (ii)*Average Delivery Latency (Delay)*: Packet delivery latency averaged over all multicast packets delivered to all receivers. It includes all possible delays caused by queuing at the interface queues, backoff at the MAC layer when the channel is busy, as well as propagation and transfer times. (iii)*Forwarding Cost (FC)*: The total number of data packet transmissions divided by the total number of packets received by all the multicast members. It gives the average number of transmissions required per delivered packet. In an ideal environment, the number of data received (denominator) is same for all protocols, and hence this metric degenerates to be the same as the total number of transmissions. In a realistic environment, the PDR is different for each protocol, and hence this metric along with the total number of transmissions gives a better picture of the forwarding efficiency of each protocol.

4.2. Results

Figure 5 shows the performance of the three protocols under different network sizes. The main observations are as follows.

In Figure 5(a) we observe that PDR is less than 100% for all the three protocols, since packet loss occurs due to bit errors and due to collisions and it drops with the network size. But when the network size changes from 500 to 1000 nodes, this drop is negligible for HRPM and HGMR (from 84% to 82%), although significant for GMR (from 72% to 60%). Similarly, in Figure 5(b) we observe that for a 1000-node the network delivery latency for GMR is much higher compared to the other two protocols (68 msec for GMR vs. 54 msec for HRPM and 53 msec for HGMR).

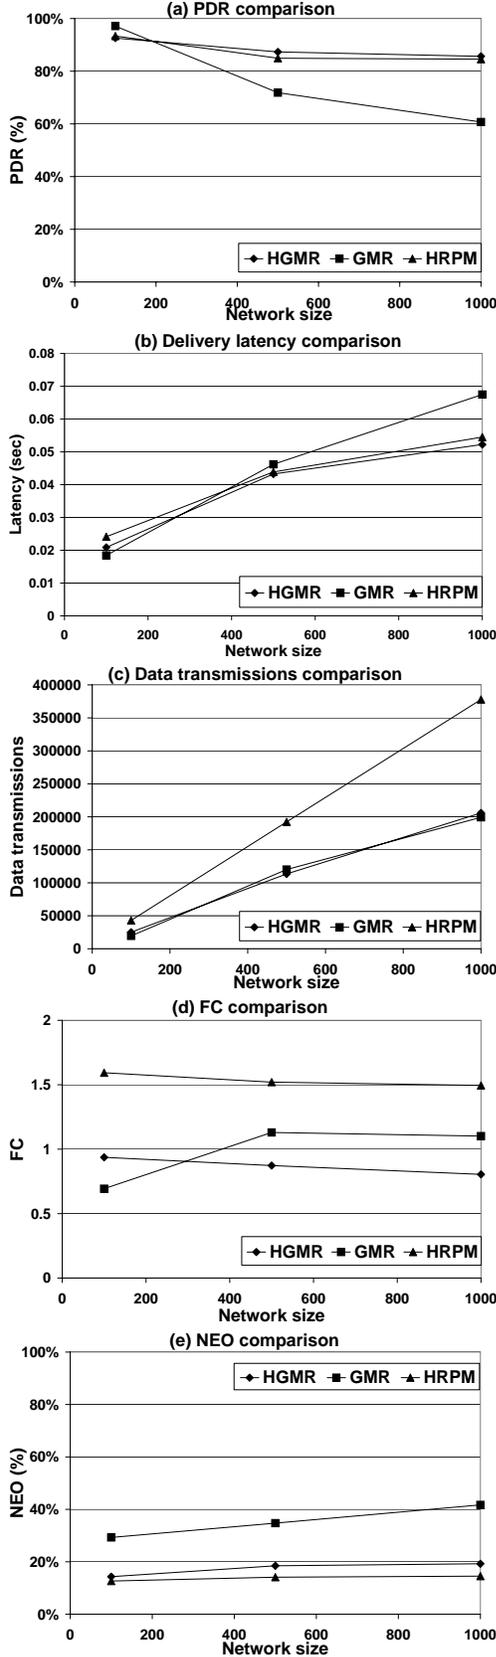


Figure 5. Performance comparison of GMR, HRP and HGMR for different network sizes in a realistic environment.

This shows that GMR does not scale well to large network sizes. On the other hand, the performance of the other two protocols, HGMR and HRP is almost identical; HGMR only slightly outperforms HRP. This difference in terms of PDR and delivery latency among the three protocols can be explained if we look at the rest of our evaluation metrics in Figures 5(c), 5(d), and 5(e).

Figure 5(c) shows that as in the ideal environment, HRP incurs a much higher number of transmissions compared to the other two protocols, and the gap increases with the network size – it reaches 88% for 1000 nodes. On the other hand, GMR and HGMR have similar numbers of transmissions. Note also that the total number of transmissions is lower for all the three protocols, compared to the ideal case in Figure 4(a), due to packet losses. We have also seen in Figure 4(b) that HRP and GMR use many more nodes as forwarders compared to HGMR – the number of FNs remains the same under ideal or realistic conditions.

Figure 5(d) shows that the forwarding cost drops with the network size for HGMR and HRP, but increases for GMR. FC combines the number of transmissions and the number of packets received by all members, and hence it gives a better idea of the forwarding efficiency of a protocol in a realistic environment. HGMR is the most efficient protocol, maintaining a FC lower than 1 for all network sizes, due to the WMA. This means that a single transmission with HGMR is enough for more than one receivers. On the other hand, HRP is the costliest protocol, with FC always larger than 1.5. Interestingly enough, GMR does not follow the trend of the other two protocols. Its FC increases with the network size, although it also exploits WMA to a higher degree than HGMR. However, as the network size increases to 500 or more nodes, PDR drops significantly for GMR, as Figure 5(a) shows, and this reduction in PDR affects FC, which becomes larger than 1.

This low performance of GMR for large networks is due to the large packet size. In Figure 5(e) we observe that NEO for GMR becomes larger than 40% in a 1000-node network, while it is only 16% for HGMR and 14% for HRP. An increase in the packet size increases both the probability of bit errors (which causes packet corruption) and the probability of packet collisions. The later is also exacerbated by the large number of FNs (in Figure 4(b) we saw that the percentage of FNs increases with the network size for GMR while it remains constant for the other two protocols). This results in packet loss which reduces the PDR. Also, contention among neighboring nodes increases, and this increases packet delivery latency. For HRP, the probability of collisions should also increase, in this case due to the large number of transmissions. But the low NEO in Figure 5(e) shows that the packet size is kept small for HRP, hence the probability of collisions finally remains low, as well as the probability of bit errors. For this reason HRP maintains both high PDR and low delivery latency

for all packet sizes, in spite of being wasteful in terms of network resources. Finally HGMR maintains a much lower number of transmissions compared to HRPm and a much lower NEO compared to GMR, and by balancing these two factors, it can achieve the highest performance and the best network resource conservation among the three protocols.

5. Related Work

In addition to GMR and HRPm, HGMR is related to previous location-based multicast protocols and hierarchical non-location-based multicast protocols. Early location-based protocols [4, 5, 6] were proposed for small groups due to the constraint of encoding either the entire tree or the destinations in the data packet headers. The SPBM protocol [14] shares with HGMR and HRPm the essence of improving the scalability of location-based multicast using hierarchical group management. However, SPBM uses flooding in hierarchical group management, while HGMR/HRPm use mobile geographic hashing. In addition to GMR [1], several other geographic multicast protocols target different ways of constructing multicast trees in conjunction with geographic routing, either in a general setting or for sensor networking [15, 16, 17].

Several hierarchical non-location-based protocols have been proposed which can be overlay or non-overlay based. Protocols such as AMRIS [18] and PAST-DM [19] propose an overlay-based approach in which the overlays are a form of hierarchies. An example of a non-overlay hierarchical multicast protocol is HDDM [20] which extends DDM to include a hierarchical structure. Finally, the work in [21] proposed the use of *cores* to reduce control traffic for creating multicast delivery structures.

6. Conclusions

In this paper, we have presented HGMR, a new location-based multicast protocol for wireless sensor networks. HGMR seamlessly incorporates innovations in location-based multicast and optimizes them for wireless sensor networks by providing both energy-efficiency as well as scalability to large networks. Our simulation studies confirm that HGMR combines the strengths of the two protocols it leverages (HRPm and GMR): In an ideal environment, HGMR incurs a number of transmissions either very close to or lower than GMR, and an encoding overhead very close to HRPm, as the group size increases with the network size; in a realistic environment, HGMR achieves a PDR close to HRPm and much higher than GMR, and the lowest packet delivery latency among the three protocols, while incurring much fewer packet transmissions than HRPm. For future work, we plan to extend HGMR to support different cost functions, as well as manycast and anycast services.

Acknowledgment

This work is partially supported in part by NSF grant ANI-0338856, NSERC Strategic Grant STPGP 336406, and the UK Royal Society Wolfson Research Merit Award.

References

- [1] J. Sanchez, P. Ruiz, X. Liu, and I. Stojmenovic, "GMR: Geographic Multicast Routing for Wireless Sensor Networks," in *Proc. of IEEE SECON*, 2006.
- [2] S. M. Das, H. Pucha, and Y. C. Hu, "Distributed hashing for scalable multicast in wireless ad hoc networks," *IEEE TPDS*, 2007.
- [3] USCG Navigation Center GPS page, "http://www.navcen.uscg.ni/gps/default.html," January 2000. [Online]. Available: <http://www.navcen.uscg.ni/gps/default.html>
- [4] S. Basagni, I. Chlamtac, and V. Syrotiuk, "Location aware, dependable multicast for mobile ad hoc networks," *Computer Networks*, vol. 36, pp. 659–670, August 2001.
- [5] K. Chen and K. Nahrstedt, "Effective Location-Guided Tree Construction Algorithms for Small Group Multicast in MANET," in *Proc. of IEEE INFOCOM*, June 2002.
- [6] M. Mauve, H. Fuessler, J. Widmer, and T. Lang, "Position-based multicast routing for mobile ad-hoc networks," University of Mannheim, Tech. Rep. CS TR-03-004, 2003.
- [7] E. M. Royer and C. E. Perkins, "Multicast operation of the ad-hoc on-demand distance vector routing protocol," in *Proc. of MobiCom*, August 1999.
- [8] J. G. Jetcheva and D. B. Johnson, "Adaptive Demand-Driven Multicast Routing in Multi-Hop Wireless Ad Hoc Networks," in *Proc. of ACM MobiHoc*, October 2001.
- [9] S.-J. Lee, M. Gerla, and C.-C. Chiang, "On-Demand Multicast Routing Protocol," in *Proc. of IEEE WCNC*, September 1999.
- [10] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia, "Routing with guaranteed delivery in ad hoc wireless networks," in *Proc. of ACM DialM Workshop*, August 1999.
- [11] B. Karp and H. Kung, "GPSR: Greedy perimeter stateless routing for wireless networks," in *Proc. of ACM MobiCom*, August 2000.
- [12] F. Kuhn, R. Wattenhofer, and A. Zollinger, "Worst-case optimal and average-case efficient geometric ad-hoc routing," in *Proc. of ACM MobiHoc*, June 2003.
- [13] X. Zeng, R. Bagrodia, and M. Gerla, "Glomosim: A library for parallel simulation of large-scale wireless networks," in *Proc. of PADS Workshop*, May 1998.
- [14] M. Transier, H. Fuessler, J. Widmer, M. Mauve, and W. Effelsberg, "Scalable Position-Based Multicast for Mobile Ad-hoc Networks," in *Proc. of the First International Workshop on Broadband Wireless Multimedia: Algorithms, Architectures and Applications (BroadWin)*, October 2004.
- [15] P. M. Ruiz and I. Stojmenovic, "Cost-efficient Multicast Routing in Ad Hoc and Sensor Networks," In *Handbook on Approximation Algorithms and Metaheuristics*, Chapman and Hall/CRC Teofilo Gonzalez, ed., 2006.
- [16] I. Stojmenovic, "Localized network layer protocols in sensor networks based on optimizing cost over progress ratio," *IEEE Network*, vol. 20, no. 1, January/February 2006.
- [17] S. Wu and K. S. Candan, "GMP: Distributed Geographic Multicast Routing in Wireless Sensor Networks," in *Proc. of IEEE ICDCS*, 2006.
- [18] C. Wu and Y. Tay, "AMRIS: A Multicast Protocol for Ad hoc Wireless Networks," in *Proc. of MILCOM*, November 1999.
- [19] C. Gui and P. Mohapatra, "Efficient overlay multicast for mobile ad hoc networks," in *Proc. of IEEE WCNC*, March 2003.
- [20] —, "Scalable multicasting for mobile ad hoc networks," in *Proc. of IEEE INFOCOM*, March 2004.
- [21] E. Madruga and J. Garcia-Luna-Aceves, "Scalable Multicasting: The Core Assisted Mesh Protocol," *ACM/Baltzer Mobile Networks and Applications, Special Issue on Management of Mobility in Distributed Systems*, vol. 6, no. 1, 2001.