Distributed Hashing for Scalable Multicast in Wireless Ad Hoc Networks

Saumitra M. Das, Himabindu Pucha and Y. Charlie Hu
Purdue University, West Lafayette, IN 47907
Email: {smdas, hpucha,ychu}@purdue.edu

Several multicast protocols for MANETs have been proposed (DSM, PBM and LGT) that build multicast trees using location information available from GPS or localization algorithms and use geographic forwarding to forward packets down the multicast trees. These *stateless* multicast protocols carry encoded membership, location and tree information in each packet. Stateless protocols are more efficient and robust than stateful protocols (ADMR, ODMRP) as they avoid the difficulty of maintaining distributed states in the presence of frequent topology changes in MANETs. However, stateless location-based multicast protocols are not scalable to large groups because they encode group membership in the header of each data packet.

In this work, we improve the scalability of stateless location-based multicast with increased group size. We design the Hierarchical Rendezvous Point Multicast (HRPM) protocol which constructs a hierarchy to organize large groups into many manageable-sized groups. More importantly, HRPM constructs and maintains this hierarchy at virtually no cost using distributed hashing. Thus, HRPM incorporates two key design concepts: (1) Hierarchical decomposition of multicast groups, and (2) Leveraging mobile distributed geographic hashing to construct and maintain such a hierarchy efficiently. We find that this architecture also allows HRPM to scale for large networks and large number of groups.

The per-packet encoding overhead of a stateless location-based multicast protocol grows with the group size G as $O(G)$. An increase in G severely limits the usability of such protocols. HRPM limits the per-packet overhead to an application-specified constant (ω) 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 MANET region into smaller and smaller *cells*. As Figure 1 shows, such cells 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. When the group size changes, HRPM dynamically adjusts the hierarchy to meet this ω.

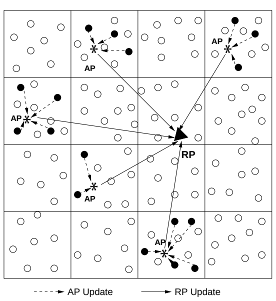

(1) Hierarchy and tree construction

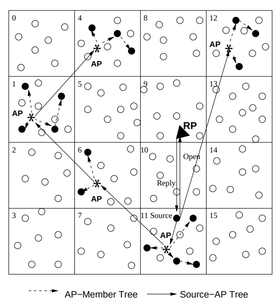

(2) Data delivery

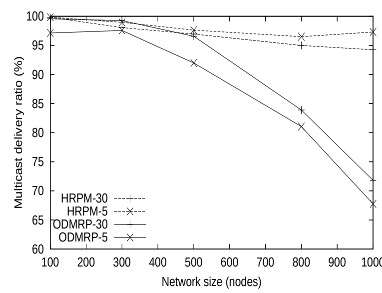

(3) Scalability

Central to the design of HRPM is the fact that both RPs and APs are logical entities. If such a logical entity is associated with a specific node (IP address), keeping track of the RP/AP would require a high overhead due to mobility in MANETs. To avoid such overhead, HRPM disassociates the RP/AP from any specific node by adopting the concept of *geographic hashing*. Given a data item and a location, geographic hashing maps (routes) the data item to the node whose geographic location is closest to the location. Since in MANETs different mobile nodes can become the closest to a fixed location over time, *mobile geographic hashing* in HRPM extends geographic hashing via a continuous handoff process which ensures the data item is always stored on the node currently closest to the location. Thus if the members of a group/subgroup use an agreed-upon hashing function to hash the multicast group identifier and obtain the RP/AP location for the group/subgroup, all group management messages can be routed to the "stationary" RP/AP by leveraging geographic forwarding. HRPM provides a framework for scalable group management in location-based multicast in which any tree construction algorithm of choice can be utilized based on the application metrics. In our study, HRPM uses a Steiner tree heuristic to optimize bandwidth. Geographic forwarding is used for data delivery between overlay tree nodes as shown in Figure 2.

We implemented HRPM in the widely used Glomosim simulator. Our simulation results show HRPM maintains can support more than 250 group members in a 500-node network. In this case, previous location-based multicast protocols fail miserably and cannot even deliver packets due to the large packet sizes. Figure 3 also shows that HRPM outperforms the state-of-the-art non-location based protocol, ODMRP, by delivering significantly more data packets as the network size is increased. For each network size, we consider group sizes of 5% and 30% of nodes in the network and we find that HRPM outperforms ODMRP regardless of the group size. In our experiments, HRPM was also found to scale well with the number of groups. In conclusion, HRPM significantly improves the scalability of location-based multicast for MANETs.