Flooding for Reliable Multicast in Multi-Hop Ad Hoc Networks*

Christopher Ho, Katia Obraczka, Gene Tsudik, Kumar Viswanath
USC Information Sciences Institute
4676 Admiralty Way, Marina del Rey, CA 90292, USA
E-mail contact: katia@isi.edu
Tel: (310) 822-1511 Fax: (310) 823-6714

Abstract

Ad Hoc Networks are gaining popularity as a result of advances in smaller, more versatile and powerful mobile computing devices. The distinguishing feature of these networks is the universal mobility of all hosts. This requires re-engineering of basic network services including reliable multicast communication.

This paper considers the special case of highly mobile fast-moving ad hoc networks and argues that, for such networks, traditional multicast approaches are not appropriate. Flooding is suggested as a possible alternative for reliable multicast and simulation results are used to illustrate its effects. The experimental results also demonstrate a rather interesting outcome that even flooding is insufficient for reliable multicast in ad hoc networks when mobility is very high. Some alternative, more persistent variations of flooding are sketched out.

1 Introduction

Recent advances in portable computing devices and wireless communication technology have made it possible to *stay connected anywhere, anytime.* In the near future, users will be able to move freely and still have seamless, reliable and high-speed network connectivity. Portable computers and hand-held devices will do for data communication what cellular phones are now doing for voice communication.

Traditional network mobility focused on roaming, which is characterized by hosts connecting to the fixed-infrastructure internet at locations other than their well-known *home* network address. Hosts can connect directly to the fixed infrastructure on a visited subnet through a wireless link or a dial-up line. These so-called traditional (or fixed-infrastructure mobile) networks raise issues such as address management but do not require significant changes to core network functions such as routing.

*This work is supported by NSF-NGI grant number ANI-9813724.

Multi-hop *Ad Hoc* Networks (AHNs) refer to – mostly wireless – networks where all network components are mobile. In an AHN there is no distinction between a host and a router since all network hosts can be endpoints as well as forwarders of traffic. In contrast with fixed-infrastructure networks, AHNs require fundamental changes to network routing protocols, including multicast routing and packet forwarding.

Little has been done to-date as far as providing support for multicast communication in AHNs. A major challenge lies in achieving reliable multicast communication in environments with universal mobility and frequent node outages and failures.

To this end, this paper explores the limits of reliable multicast in very dynamic, high-mobility AHNs and motivates the need for new multicast routing protocols aimed specifically at such networks. We focus on studying flooding as an alternative to multicast routing in very dynamic AHNs. Flooding emphasizes minimal state and high reliability which makes it very attractive for highly dynamic, fast-moving AHNs. We present preliminary simulation results that test the limits of flooding's reliability as a function of mobility.

2 Multicast in AHNs

Regardless of the underlying network environment, network-level multicast is the fundamental enabling technology for collaborative and, more generally, group communication applications. In both military (e.g., command and control in battlefield scenarios) and civilian (e.g., disaster relief) environments, group-oriented services such as teleconferencing and data distribution are anticipated to be some of the key applications for AHNs.

The challenges raised by multicast routing and packet forwarding in AHNs are essentially due to their unconstrained mobility characteristics. Unconstrained mobility implies the following:

- Individual host behavior independent of other hosts.
- Essentially no limit on host speed.
- No constraints on direction of movement.
- High probability of frequent, temporary network partitions.

The above translates into frequent topology changes which makes it difficult for a host to maintain timely multicast-related state information (membership) other than its own. Furthermore, in many types of AHNs (e.g., where hosts are hand-held devices) both storage capacity and power are severely limited. This is yet another reason to avoid maintaining and exchanging multicast state.

Another important consideration has to do with the mission of AHNs. In the critical environments AHNs are most often deployed (both in military and other emergency situations), **robustness** and **high quality-of-service** are of paramount concern. Thus, multicast mechanisms (however attractive otherwise) that cannot provide the highest delivery guarantees may not be appropriate.

Only recently, routing protocols for AHN multicast have become an active area of research. Examples of current research efforts include the On-Demand Multicast Routing Protocol (ODMRP) from UCLA [8], the multicast extensions to CEDAR by UIUC [14], and the Ad hoc Multicast Routing protocol utilizing Increasing id-numberS (AMRIS) from National University of Singapore and GeorgiaTech [16]. The common denominator among these approaches is that they all require state to be maintained by network elements. In contrast, flooding results in minimal state retention and high reliability which makes it a viable candidate for multicast protocols in very dynamic AHNs.

We also note that a single cure-all multicast solution for all AHNs is highly unlikely. Relatively stable AHNs with few host failures, outages and infrequent movements will lend themselves to approaches different from those best-suited for highly dynamic and highly volatile AHNs. Consequently, one of the long-term challenges is in determining the best multicast approach when faced with a specific AHN configuration and parameters. We envisage a suite of solutions, each geared towards (and nearly optimal in the context of) a specific AHN type and host mobility pattern. This is very much in line with multicast development in fixed networks such as the Internet. The two modes of Protocol-Independent Multicast (PIM) [12] (sparse and dense modes) are a case in point.

3 Desired Multicast Properties

Multicast routing in highly-mobile AHNs must emphasize the following:

Robustness versus Efficiency. Many multicast routing approaches rely on state in routers to keep track of multicast group members. This, coupled with the high volume of routing information exchanges and slow convergence make traditional multicast approaches untenable in highly dynamic AHNs composed of anemic (low-power, low storage capacity) hosts. Therefore, new techniques that stress rapid and robust delivery must be developed.

Adaptability. AHN behavior can change over time, i.e., a very mobile AHN can stabilize in part or as a whole, or, similarly, a relatively static AHN can suddenly become very mobile. Different multicast mechanisms are appropriate for high-mobility and low-mobility AHNs. (The same holds for sparse and dense AHNs.) Ideally, hosts should be able to adapt to AHN behavior changes by dynamically switching among different multicast mechanisms. This must be done with the minimum of both effort and inconvenience (e.g., packet loss). Also, over time, an AHN can experience drastic changes in terms of its collective mobility or, conversely, its stability.

Unlimited Mobility. Some multicast solutions are geared towards *discrete* mobility whereby periods of movement are interspersed with periods of rest. Some others assume limits on direction, speed and number of simultaneously moving hosts. In contrast, we stress *continuous* and high mobility of *all* AHN components.

Integrated Multicast. Multicast solutions for AHNs will most likely differ substantially from those for fixed networks (one of the main reasons is the marked difference in transmission rates). In order to offer seamless and integrated multicast service, new mechanisms must be developed for inter-operation of fixed and wireless multicast solutions.

4 Problem Scope

As mentioned above, our focus is on multicast routing in highly-mobile AHNs. This is a challenging and, until recently, relatively little-explored area. We note that, in contrast, multicast routing in fixed-infrastructure mobile networks presents only a few engineering obstacles. These issues (which mainly have to do with the last hop delivery and membership tracking) are not discussed in this paper.

Multicast mechanisms for other, less dynamic types of AHNs have been proposed. As discussed in Section 8, relatively slow-moving AHNs lend themselves to adaptations of traditional, state-based multicast methods such as PIM [12, 5]. Alternatively, slow-moving AHNs can be amenable to multicast extensions of discovery-based (or on-demand) routing methods such as DSR [9][1]. Also, multicast in hierarchical or clustered AHNs with fairly static cluster-level topology and little inter-cluster migration can be accommodated by traditional multicast methods for inter-cluster routing [13].

Our starting hypothesis is that reliable multicast in very dynamic high-speed AHNs is untenable with traditional multicast mechanisms, which are based on state built by routers. On-demand AHN multicast methods operate by amortizing the cost of flooding-based destination discovery over subsequent data packets. However, collecting paths to a set of destinations is useless if it is all but certain that the destinations will move by the time the next packet is sent. Proactive state-based methods derived from PIM (e.g., ST-WIM) suffer from the same problem. Keeping accurate state about the multicast group membership of all nodes' neighbors is difficult if the set of neighbors changes at a very high rate.

In light of the above it seems natural to consider plain flooding which, although heavy-handed in terms of overhead, typically obtains the best results in terms of reliable delivery. However, we postulate that – in the types of AHNs we consider – flooding does not provide the high reliability comparable to that it provides in static networks or more stable AHNs. High-speed node movement can preclude reception of a packet.

Simulations allow us to demonstrate this hypothesis by deriving concrete AHN scenarios where plain flooding fails. Informally, *fails* means the following:

> Assume that p represents per hop packet transmission delay and d is the maximum AHN diameter. A packet is sent at time t_0 and flooded throughout the AHN. Every node that receives the packet broadcasts it (exactly once) to all immediate neighbors. Let $t_1 \leq t_0 + dp$. Assuming that the AHN stays continuously connected[2] between time t_0 and t_1, there exists at least one node that has not received the packet.

5 Parameters

Simulations are invaluable when evaluating the performance of network protocols. They allow extensive exploration of a protocol's design space. This includes subjecting the protocol to extreme or boundary conditions which are often hard to reproduce in real-life (or live) experiments.

Thus, when simulating a protocol, it is essential to determine both the metrics by which to evaluate the protocol's performance, as well as the dimensions of the protocol's design space. The latter are typically used as parameters of the simulator representing knobs that can be independently turned in order to subject the protocol to a range of conditions.

This exercise is particularly hard when evaluating multicast protocols for AHNs. AHNs add **movement**, a complex dimension, to the already intricate multicast protocol design space. Since AHNs are a relatively incipient research area, well-established mobility patterns to be used when simulating AHNs are not currently available.

Aside from node mobility, the other design space dimensions of multicast routing in AHNs are listed below. We note that this is not an exhaustive parameter list.

- **Application-specific** parameters include total number of messages sent by each node or the node's transmit rate, and the interval nodes wait before forwarding a packet they receive.
- **Multicast group** parameters include number of nodes or node density. Since we are simulating broadcast, the total number of nodes is equal to the number of recipients. In the case of multicast, these values are typically different. We set the number of nodes high enough to be "interesting", but again low enough to readily simulate. We specifically avoided a number so low as to allow long-lived network partitioning.
- **Terrain-related** parameters include the dimensions of the field perimeter. In the case of three-dimensional terrains, one can also specify height. This is useful when simulating physical obstacles; however mobility-capable network simulators like GloMoSim or *ns-2* with mobility extensions by CMU [6] currently do not support this.
- **Node capability** parameters include the node's transmission power range and bandwidth.

6 Simulation Environment

For the simulations performed in this study, we used the Global Mobile System Simulator (GloMoSim) developed at UCLA. GloMoSim is a library-based sequential and parallel simulator for wireless networks [17]. It is designed as a set of library modules, each of which simulates a specific wireless communication protocol in the protocol stack. The library has been developed using PARSEC, a C-based parallel simulation language [1],

[1] A multicast-flavored DSR would use per-source flooding to build a path tree to multicast group members.

[2] Connected means that there is always a path between any two nodes.

which can be used to program new protocols and modules that can be added to the library. GloMoSim has been designed to be extensible and composable. The communication protocol stack for wireless networks is designed using a layered approach, where each layer has with its own API. Models of protocols of one layer interact with those of other layers via these APIs.

6.1 Modifications to GloMoSim

For our simulations, we used GloMoSim v1.1.1. This version has two mobility models,*random waypoint* and *drunken mobility*. In the random waypoint model the nodes start off with random positions in the terrain. Each node randomly selects a destination and moves in the direction of the destination, moving one meter every *mobility-interval* time period. After it reaches its destination, the node stops for *mobility-pause* time period and then chooses a new destination. In the *drunken mobility model* each mobile node is assigned a random position within the field. When the node is next considered for movement, the mobility module checks all the possible directions in which the node can move to ensure that it stays within the field boundaries. The node then moves in the direction randomly chosen from the set of possible directions. Each node moves by one unit distance in that direction during the *mobility-interval*.This movement pattern is also known as *random walk*.

When we traced node movement for the *drunken mobility* model during the simulation, we found that nodes generally exhibited oscillatory movement patterns, i.e., they tended to move back and forth about their original positions in the field. Consequently, nodes never moved significantly in any direction.

To overcome this problem we modified GloMoSim's original mobility model in order to bias node movement towards the direction randomly chosen when the node is first selected for movement. When the node reaches the field boundary, it chooses and proceeds in a new direction. This model is essentially the *random-waypoint* model with the destinations set to the terrain boundaries and the *pause-time* set to zero seconds.

Parameter	Value	Description
number-of-nodes	50	simulation nodes
num-packets	25	messages sent by a node
flooding-interval	10 ms	random wait before flooding
field-range-x	1000 m	X-dimension of motion
field-range-y	1000 m	Y-dimension of motion
power-range	250 m	node's power range
bandwidth	2 Mbit/s	node's bandwidth
mobility-interval	10-100 ms	
distance-unit	1m	
simulation-time	100 s	simulation duration
node-placement	random	node placement policy
propagation-func	FREE-SPACE	propagation function
radio-type	NO-CAPTURE	capture effect
mac-protocol	CSMA protocol	MAC layer
network-protocol	flooding	network layer
transport-protocol	UDP protocol	transport layer

Table 1: Simulation Parameters.

6.2 Flooding Application

Our flooding application simultaneously runs on n nodes which are randomly placed in the simulated field. Throughout the course of the simulation, each of the nodes attempts to broadcast m messages to all other nodes in the network. Thus the application expects the nodes to collectively receive $m*(n-1)$ messages. When a node receives a packet, it waits a uniformly distributed time interval between 0 and *flooding-interval* before it broadcasts the packet. n, m, and *flooding-interval* are parameters of the simulator.

We count each time a message from any sender fails to reach any receiver. The difference between the expected and the actual number of messages collectively received divided by the total number of packets sent is the *packet loss factor*. Note that we are actually performing broadcast, a special case of multicast.

6.3 Simulation Platform, Parameters, Methodology

We ran the simulations on a Sun Ultra 5 with 128 megabytes of memory running Solaris 2.

Table 1 summarizes the simulation parameters we used. The total simulation time was 100 seconds. A total of 50 nodes were randomly placed in a field of a 1000x1000 meters. Each node generated 25 messages using a 2Mbit/sec channel with power range of 250 meters. The rationale for using this power range is explained in Section 7 below. In our simulations, *flooding-interval* is set to *10 ms*. Free space propagation was used to determine whether nodes that are in the transmitter's range are able to receive data. The protocol stack we used consisted of CSMA, flooding and UDP as the MAC, network and transport layer protocols, respectively.

In our simulations each node transmits 25 packets over a period of 75 seconds (3/4 of the total simulation time). Since all nodes are sources, this translates to 1250 packets in 75 secs or a network traffic load of 17 packets/sec. We should also point out that generating traffic within the first three quarters of the simulation ensures that all nodes have a chance to receive all packets before the simulation ends.

We ran each simulation (maintaining all simulation parameters the same) ten times, each time with a different seed value. Seeds varied from 1000 to 10000 in steps of 1000. The graphs presented in Section 7 below plot the average across all ten runs.

7 Results

7.1 Setting the Power Range

A critical factor influencing packet loss in AHNs is the effective transmission range of each node. We first studied this parameter to get a better understanding of its effect and also to establish optimum values which could be used for our simulations. Figure 1 shows the results obtained.

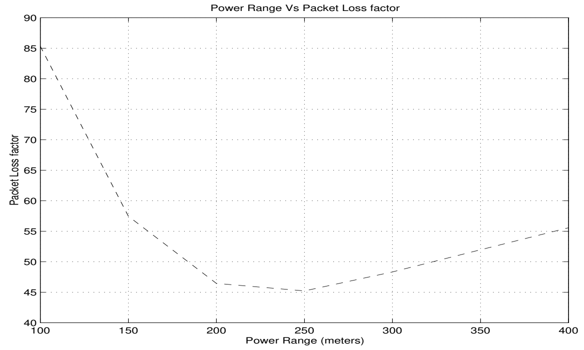

Figure 1: Packet Loss as a function of the Power Range

The simulations were performed for 50 nodes randomly placed in a 1000x1000 field, using 7 different power ranges from 100 to 400 meters. The *mobility interval* for all the simulations was set to 100 milliseconds (36 km/hr). Each node transmitted 25 messages at random times during the simulations. At each power the test was run 10 times using 10 different random number seeds. From Figure 1 it can be observed that the "sweet spot" in the curve is in the range from 200 to 250 meters. This is due to the fact that at low power ranges, packet losses are high largely due to node disconnectivity. In the other extreme, i.e., for higher power ranges, losses are mostly due to collisions. For our simulations we set the transmission range for each node to be 250 meters.

7.2 The Effects of Mobility

Our goal was to study how node mobility affects the performance of flooding in wireless networks. More specifically, we were interested in evaluating the effects of mobility on flooding's ability to deliver packets reliably. We use the packet loss factor (as defined in Section 6) to measure the percentage of packets lost.

Recall that flooding works by having each node forward a packet it receives out on every link except the one from which it received the packet. In wired networks, flooded packets can be lost due to network partitions (caused by link or node failure) that happen ahead of the "flooding wave" and are not mended in time. In AHN environments, a node may miss a packet if it moves out of range as the flooding wave goes by. Unless the node subsequently migrates to a region of the network where the wave has not gone by, it will miss the packet.

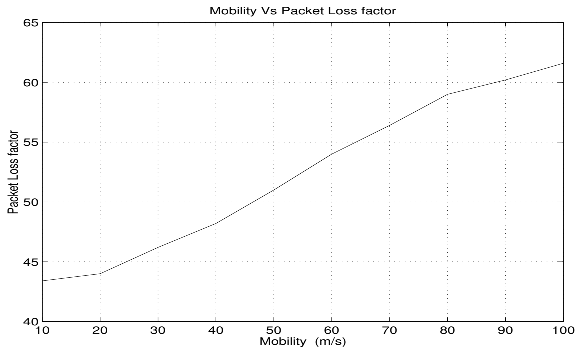

Figure 2: Packet Loss as a function of Mobility

In our simulations we varied the node speeds from no mobility to 100 m/s in steps of 10 m/s. Each experiment was run ten times with ten different random number seeds and the packet loss factor was then averaged over the entire set for a given speed. Figure 2 shows flooding's packet loss factor as a function of node mobility.

From the figure it can be observed that the packet loss factor increases from around 44% at node speeds of 10 m/s (36 km/hr) to around 62% at speeds of 100 m/s (360 km/hr) . The nodes' high speeds cause them to move out of receiving range from any of their neighbors resulting in higher packet loss. It should be noted that in sparse groups, even low mobility can cause nodes to move out of range resulting in much higher loss values.

As previously mentioned the network traffic load was approximately 17 packets/sec. More recent results show that, due to collisions, increasing network traffic results in a higher packet loss even at low speeds.

7.3 Additional Observations

Amount of overhead generated is an important criterion when evaluating the performance of communication protocols. Since flooding maintains minimal amount of state information, its overhead is primarily due to multiple copies of the same packet circulating in the network.

Figure 3 shows the number of duplicate packets received by each node as a function of mobility. The average number of duplicate packets received by a node is seen to decrease as the mobility increases. In our simulations with 50 nodes in a 1000x1000 field, each node transmits 25 messages which means that each node should ideally receive 1225 unique messages. From the figure it can be observed that each node receives 2623 duplicate messages at a node speed of 10 m/s and this reduces to 1864 messages at a speed of 100 m/s. Thus, on average, a node receives about 2.1 duplicates for each transmitted packet when the mobility is 10 m/s and this decreases to 1.5 packets when the mobility is 100 m/s. We hypothesize that this increase in the number of duplicate packets received by a node with decreased mobility is due to the fact that slower moving nodes capture higher percentage of transmitted packets.

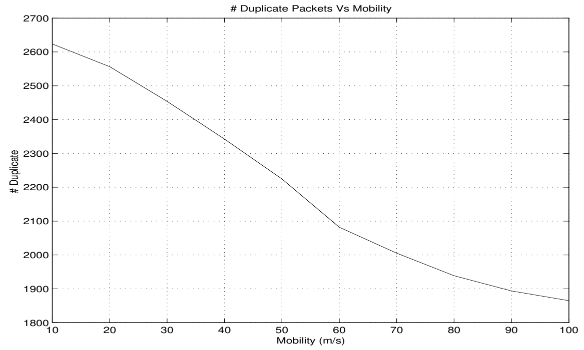

Figure 3: Duplicate packets as a function of mobility.

One inherent problem with flooding is that the broadcast nature of the protocol results in large number of collisions [10]. In a CSMA network since there is no RTS/CTS mechanism and in the absence of collision detection, collisions are more likely to occur. Figure 4 shows the number of collisions as a function of the mobility. From the figure it can be seen that on an average there are around 2667 collisions at each node when the speed is 10 m/s and it reduces to 1266 at node speeds of 100 m/s. At low mobility the nodes are within transmission range of a greater number of nodes which results in larger number of collisions but as the nodes begin to move faster they are within transmission range of a significantly smaller number of neighbors resulting in lower number of collisions.

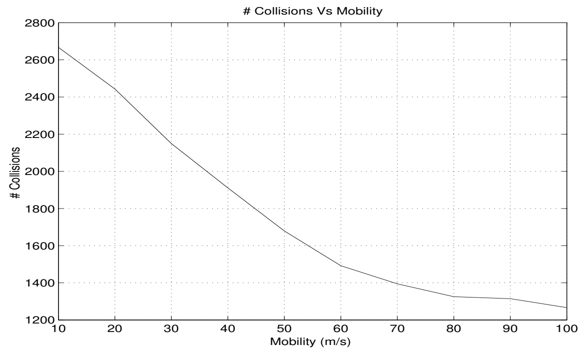

Figure 4: Collisions as a function of mobility.

8 Related Work

There are a number of proposed multicast protocols oriented towards AHN environments. These can be grouped in two main categories: proactive protocols that maintain routing state, and reactive protocols that acquire routes on demand. In this section we briefly describe several protocols in both categories.

8.1 Proactive Protocols

The Ad Hoc Multicast Routing Protocol (AMRoute) [3] builds a bidirectional shared tree to multicast data to member nodes. Communication between these nodes is layered on conventional unicast messaging. *Core nodes* are responsible for member detection and group setup, but unlike *core-based trees* [2], these core nodes do not pass data, and can migrate dynamically in response to membership and connectivity changes. AMRoute relies on the underlying unicast protocol to handle topology changes, with a tunneling mechanism used for traffic between "multicast islands" traversing regions where multicast is not deployed. This is advantageous in that intermediate routers need not run any multicast protocol, and overhead is confined to nodes participating in the multicast groups. The penalty for this user-level approach is poorer efficiency in multicast packet replication and in packet delivery delay.

Although directed acyclic graphs (tree structures) are conventionally used for efficient multicast on wired networks, the Core-Assisted Mesh Protocol (CAMP) [7] proposes *multicast meshes* for improved robustness on mobile networks. Topology changes do not necessarily trigger multicast reconfiguration. Extending the idea of core-based trees, CAMP builds on a traditional architecture for wired multicast, but uses reverse shortest paths (the shortest path from receiver to source) instead of a top-down tree. This improves efficiency over the shared multicast tree approach. *Heartbeats* are used to maintain the reverse shortest paths. In contrast to FGMP [4], CAMP specifically avoids the need to flood the network with either data or control messages, which is seen as unscalable.

The MCEDAR protocol [15] is layered as a multicast extension on top of Core-Extraction Distributed Ad Hoc Routing (CEDAR) [14]. In CEDAR, a set of hosts are selected as the *core* of the network, approximating a *minimum dominating set*. Every host within this core is responsible for route computation for its local mobile hosts not in the core. Long-lived links are preferred when building the core graph; thus local topology changes infrequently cause global updates. Link-up transitions are distributed slowly but link-down conditions are propagated quickly. A *core broadcast* mechanism is implemented using reliable unicast messages that scales linearly with the number of network nodes. CEDAR emphasizes quality-of-service (QoS) requirements.

Like CAMP, MCEDAR uses the idea of multicast meshes, adopting the notation of k-connectivity, where k denotes the maximum number of logical connections at each multicast member. However, multicast traffic is distributed using only a spanning tree subgraph of the k-connected mesh. Thus MCEDAR claims the advantages of both multicast trees – efficiency – and multicast meshes – reliability. Join operations travel through the core until they reach a core node in the appropriate multicast mesh, and are accepted up to the k connectivity factor. Leave operations update the multicast mesh appropriately and may be deliberate or caused by network partition.

The Ad Hoc Multicast Routing protocol utilizing Increasing id-numberS (AMRIS) [16] builds a shared delivery tree for the multicast group. It assigns an identification number to each multicast member in the group; the root of the tree has the lowest number in the tree, with numbers increasing radially outwards. Recovery from broken links is done locally within one period of neighbor discovery, with route reconstruction favored over route discovery. This is accomplished via an expanding ring search by the higher numbered node. Recall the higher numbered node is further from the root, so this is similar to a receiver-initiated join request. Like CEDAR, long-lived links are preferred when building the tree, to reduce later reconfiguration.

8.2 On-Demand Protocols

The Ad Hoc On Demand Distance Vector (AODV) routing protocol [11] proposes both a unicast and multicast mechanism for data delivery. In contrast to the previous proactive approaches, the on-demand trait means routes are only created for active traffic, and thus overhead is minimized for routing table space, processing, and transmission bandwidth. AODV resembles traditional link-state or distance-vector routing protocols, but modified for the AHN environment to send fewer messages. Route request messages are broadcast in an expanding ring to find a path to a destination node; route replies are unicast back. To guard against the "counting to infinity" convergence problem, AODV stamps routing information with sequence numbers; later sequence numbers are always preferred. AODV implements neighbor discovery using IP datagrams, which cause overhead. AODV performs multicast by maintaining a multicast tree for each multicast group. The tree members can be multicast group members or intermediate routers connecting the group members. Each node maintains a routing table entry for each multicast group for which the node is a member or a router. The multicast group leader periodicially broadcasts *Group Hello* messages to maintain the group connectivity. A node initiates a *RREQ* packet when it wishes to join a multicast group. The nodes receiving the *RREQ* packet set up reverse paths to the source in their unicast routing table. When a group member receives the *RREQ* it responds with a *RREP* packet.

The On-Demand Multicast Routing Protocol (ODMRP) [4] also uses an on-demand reactive approach based on *forwarding groups*. Although not as formally structured as meshes or cores, FGMP uses a subset of the connected nodes to be responsible for forwarding multicast traffic. This is referred to as *scoped flooding*. Membership in these forwarding groups require receiver-initiated control messages.

8.3 Discussion

All the above multicast protocols maintain state, some more than others, in the form of a multicast tree or mesh, which requires overhead to construct and maintain. Even the on-demand protocols require a node to have knowledge of its multicast neighbor(s), which may be instantaneously accurate but soon stale. In scenarios involving **extremely high mobility**, shortcomings become evident because node mobility causes conventional multicast trees to rapidly become outdated. Frequent state changes require constant routing updates, possibly never converging to accurately portray the current topology. Each protocol optimizes different considerations; in situations with high mobility we propose dispensing entirely with the structure of the multicast hierarchy, and using a flooding strategy.

9 Summary and Future Work

In closing, this paper discussed a number of features and characteristics of highly mobile fast-moving AHNs and argued that high speed and frequent mobility intervals necessitate radical approaches for reliable multicast. In particular, flooding is proposed as a stateless and topology-independent mechanism for reliable multicast.

The simulation results illustrate that flooding is quite effective. However, when mobility intervals are small and node speed is sufficiently high, even flooding becomes unreliable. This is a surprising outcome which motivates further investigation of more robust and more persistent variations of flooding suitable to such volatile and dynamic AHNs.

Specifically, the following are some of the issues to be addressed:

- More robust forms of flooding would likely incur some per packet state in network nodes. Although not as long-term in nature as the state kept in tree-based multicast mechanisms, the amount of state must be kept to a minimum. Extensive experiments and heuristic adjustments are needed to shed some light on this issue.

- The overhead of flooding and its variations needs to be carefully measured. We note that the overhead of flooding in fixed networks is comparatively trivial to measure (a packet traverses a given network link at most once). In AHNs, the overhead is likely to be more elusive to measure and quantify since high mobility can result in multiple receptions of the same packet and "flooding waves" would take longer to complete due to mobility-induced packet loss.

- The broadcast nature of the protocol results in packet loss due to collisions. One technique of reducing the effect of collisions is to use scoped flooding. In scoped flooding only a subset of the nodes forward the packets reducing the number of redundant broadcasts. This also involves determining the optimum number of forwarding nodes to reduce the collisions but ensuring higher packet delivery ratios.

10 Acknowledgments

We are indebted to Jeremy Elson for some of our ideas and many of the early modifications to GloMoSim.

References

[1] R. Bagrodia and R. Meyerr. PARSEC: A parallel simulation environment for complex system. Technical report, UCLA, 1997.

[2] A. Ballardie, P. Francis, and J. Crowcroft. Core based trees (CBT): An architecture for scalable inter-domain multicast routing. In *Proc. of ACM SIGCOMM'93*, 1993.

[3] E. Bommaiah, M. Liu, A. McAuley, and R. Talpade. AM-Route: Adhoc multicast routing protocol. IETF manet (draft-talpade-manet-amroute-00.txt), Work in Progress, 1998.

[4] C. Chiang and M. Gerla. On-demand multicast in mobile wireless networks. In *Proc. IEEE ICNP '98*, 1998.

[5] C. Chiang, M. Gerla, and L. Zhang. Shared tree wireless network multicast. In *IEEE International Conference on Computer Communications and Networks (ICCCN'97)*, September 1997.

[6] CMU Monarch Project. *Mobility Extensions to ns-2*, 1999. Available from http://www.monarch.cs.cmu.edu/.

[7] J. Garcia-Luna-Aceves and E. Madruga. A multicast routing protocol for ad-hoc networks. In *Proc. of INFOCOM '99*, 1999.

[8] M. Gerla and S.J. Lee. On-demand multicast routing protocol for mobile ad-hoc networks. Available from http://www.cs.ucla.edu/NRL/wireless/.

[9] D. Johnson and D. Maltz. Dynamic source routing in ad hoc wireless networks. In Tomasz Imielinski and Hank Korth, editors, *Mobile Computing*, chapter 5. Kluwer Academic Publishers, 1996.

[10] S. Ni, Y. Tseng, Y. Chen, and J. Sheu. The broadcast storm problem in a mobile ad hoc network. In *ACM MOBICOM '99*, August 1999.

[11] C. Perkins and E. Royer. Ad hoc on demand distance vector (AODV) routing. IETF manet (draft-ietf-manet-aodv-05.txt), Work in Progress, 2000.

[12] PIM development group. *Protocol Independent Multicast*, November 1997. Available from http://netweb.usc.edu/pim/.

[13] S. Ramanathan and M. Steenstrup. A survey of routing techniques for mobile communications networks. *ACM/Baltzer MONET*, pages 89–104, October 1996.

[14] P. Sinha, R. Sivakumar, and V. Bharghavan. CEDAR: Core extraction distributed ad hoc routing. In *Proc. of IEEE INFOCOMM '99*, 1999.

[15] P. Sinha, R. Sivakumar, and V. Bharghavan. MCEDAR: Multicast core extraction distributed ad-hoc routing. In *Proc. of the Wireless Communications and Networking Conference*, 1999.

[16] C. Wu, Y. Tay, and C. Toh. Ad hoc Multicast Routing protocol utilizing Increasing id-numberS (AMRIS). IETF manet (draft-ietf-manet-amris-spec-00.txt), Work in Progress, 1998.

[17] X. Zeng, R. Bagrodia, and M. Gerla. GloMoSim: a library for parallel simulation of large-scale wireless networks. In *Proc. of PADS'98*, 1998.