Exploring Mesh and Tree-Based Multicast Routing Protocols for MANETs

Kumar Viswanath, Katia Obraczka, *Member*, *IEEE*, and Gene Tsudik

Abstract—Recently, it became apparent that group-oriented services are one of the primary application classes targeted by MANETs. As a result, several MANET-specific multicast routing protocols have been proposed. Although these protocols perform well under specific mobility scenarios, traffic loads, and network conditions, no single protocol has been shown to be optimal in all scenarios. The goal of this paper is to characterize the performance of multicast protocols over a wide range of MANET scenarios. To this end, we evaluate the performance of mesh and tree-based multicast routing schemes relative to flooding and recommend protocols most suitable for specific MANET scenarios. Based on the analysis and simulation results, we also propose two variations of flooding, *scoped flooding* and *hyper flooding*, as a means to reduce overhead and increase reliability, respectively. Another contribution of the paper is a simulation-based comparative study of the proposed flooding variations against plain flooding, mesh, and tree-based MANET routing. In our simulations, in addition to "synthetic" scenarios, we also used more realistic MANET settings, such as conferencing and emergency response.

Index Terms—Ad hoc networks, mobile computing, multicast, routing protocols, wireless.

1 Introduction

Mobile multihop ad hoc networks (MANETs) are characterized by the lack of any fixed network infrastructure. In a MANET, there is no distinction between a host and a router since all nodes can be sources as well as forwarders of traffic. Moreover, all MANET components can be mobile.

MANETs differ from traditional, fixed-infrastructure mobile networks, where mobility occurs only at the last hop. Although issues such as address management arise in the latter, core network functions (especially, routing) are not affected. In contrast, MANETs require fundamental changes to conventional routing and packet forwarding protocols for both unicast and multicast communication. Conventional routing mechanisms, which are based on routers maintaining distributed state about the network topology, were designed for wired networks and work well in fixed-infrastructure mobile networks. However, topology changes in MANETs can be very frequent, making conventional routing mechanisms both ineffective and expensive.

When it became clear that group-oriented communication is one of the key application classes in MANET environments, a number of MANET multicast routing protocols were proposed [7], [19], [6], [20], [21], [8]. These protocols can be classified according to two different criteria. The first criterion has to do with maintaining routing state and classifies routing mechanisms into two types: proactive and reactive. Proactive protocols maintain routing state, while the reactive reduce the impact of frequent topology changes by acquiring routes on demand.

The second criterion classifies protocols according to the global data structure used to forward multicast packets. Existing protocols are either tree or mesh-based. As in fixed (nonmobile) multicast routing, tree-based protocols build a tree over which multicast data is forwarded. Although bandwidth-efficient, tree-based protocols do not always offer sufficient robustness. Certain key features of MANETs, such as fast deployment, make them well-suited for critical environments (e.g., battlefield or disaster recovery) where robustness and reliability are essential. Thus, one of the main challenges for multicast routing in MANETs is the need to achieve robustness in the presence of universal mobility and frequent node outages. For this purpose, mesh-based protocols build a mesh for forwarding multicast data and, thus, address robustness and reliability requirements with path redundancy inherent to meshes.

The focus of our work is to explore the design space of multicast routing protocols in MANETs. More specifically, one of the goals of this paper is to characterize the merits of mesh and tree-based protocols for a wide range of MANET conditions and make recommendations for protocols best-suited to specific MANET settings. To this end, we conducted extensive simulations employing a wide range of mobility and traffic load conditions, as well as different multicast group characteristics (e.g., number of sources and number of receivers). Our study compares the performance of the On-Demand Multicast Routing Protocol (ODMRP) [7] as the representative of mesh-based protocols against Multicast Ad Hoc On-Demand Distance Vector (MAODV) [19] representing tree-based schemes. Both protocols belong to the reactive category. As a yardstick in our comparisons, we use flooding, arguably the simplest and oldest mesh-based routing technique. Despite the hefty overhead, it provides the best delivery guarantees for unicast, multicast,

- K. Viswanath and K. Obraczka are with the Computer Engineering Department, University of California, Santa Cruz, 1156 High St., Santa Cruz, CA 95064. E-mail: {kumarv, katia}@cse.ucsc.edu.
- G. Tsudik is with the Computer Science Department, University of California, Irvine, 458 CS Building, Irvine, CA 92697-3425. E-mail: gts@ics.uci.edu.

Manuscript received 28 Apr. 2003; revised 11 Apr. 2004; accepted 13 June 2004; published online 16 Nov. 2005.
For information on obtaining reprints of this article, please send e-mail to: tmc@computer.org, and reference IEEECS Log Number TMC-0057-0403.

1536-1233/06/$20.00 © 2006 IEEE Published by the IEEE CS, CASS, ComSoc, IES, & SPS

and broadcast in wired networks. However, in flooding, redundant broadcasts may cause serious contention and collision problems in MANETs. (Some of our preliminary simulation results can be found in [16].)

ODMRP was chosen since it has been shown to be the best performer in the comparative study reported in [14]. In fact, [14] compares the performance of ODMRP and CAMP [6] as mesh-based protocols against AMRoute [2] and AMRIS [21], representing tree-based mechanisms. The comparative performance study portion of this paper differs from [14] in a number of ways. First, we use MAODV as representative of tree-based multicast routing since it does not exhibit the limitations of AMRoute and AMRIS, both of which rely on an underlying unicast routing protocol. Additionally, AMRoute is susceptible to transient routing loops. Another distinguishing feature of our study is that it investigates a wider range of MANET scenarios, subjecting the protocols under consideration to more stringent network conditions, including higher mobility and traffic load, as well as a variety of multicast group characteristics (e.g., number of traffic sources, group size, and density). Finally, besides synthetic MANET environments, our study also considers more realistic scenarios such as conferencing and emergency response operations.

Based on these simulation results, we also explore the need for new protocols that provide high delivery guarantees with low overhead. Routing protocol overhead can be especially harmful in typical MANET scenarios where nodes are both bandwidth and energy-constrained. While flooding generates no control traffic, it involves redundant retransmissions. We examine *scoped flooding*, a variation of flooding that aims at reducing overhead inherent to plain flooding. Simulation results show that, at low mobility ranges (0-75 km/hr), scoped flooding achieves overhead savings of 20 percent compared to flooding and 15 percent compared to ODMRP. Interestingly, in "concrete scenarios," these overhead savings are obtained at better or comparable packet delivery ratios than ODMRP and MAODV. These overhead savings can prove to be crucial in energy constrained environments.

We also investigate another flavor of flooding, referred to as *hyper flooding*, for MANET scenarios where reliability is the primary issue. Through simulations, we show that hyper flooding can provide better reliability gains at high mobility (75-150 km/hr), which is obtained at the cost of an overhead increase compared to plain flooding. Mission-critical applications that require high reliability and timely delivery in the presence of fast-moving nodes (e.g., aircraft) may be willing to pay the price of higher overhead.

The rest of this paper is organized as follows: In the next section, we overview ODMRP and MAODV and briefly describe our implementation of flooding. Section 3 describes the simulation environment used, including a detailed description of the simulation parameters. In Section 4, we present simulation results comparing the performance of mesh (ODMRP and flooding) and tree-based (MAODV) multicast routing protocols under a variety of MANET scenarios, as well as a qualitative comparison of the protocols based on our results. Section 5 describes scoped and hyper flooding and Section 6 presents simulation results comparing their robustness and overhead relative to plain flooding, ODMRP, and MAODV. We present results for both synthetic as well as more concrete MANET scenarios. Section 7 describes related work efforts and, in Section 8, we present some concluding remarks as well as items for future work.

2 Mesh and Tree-Based Multicast Overview

In this section, we review the operation of mesh and tree-based multicast routing using ODMRP and MAODV as examples of mesh and tree-based protocols, respectively. We also highlight the main features of our implementation of flooding.

2.1 On Demand Multicast Routing Protocol (ODMRP)

The On-Demand Multicast Routing Protocol (ODMRP) [7] falls into the reactive protocol category since group membership and multicast routes are established and updated by the source whenever it has data to send. Unlike conventional multicast protocols, which build a multicast tree (either source-specific or shared by the group), ODMRP is mesh-based. It uses a subset of nodes, or *forwarding group*, to forward packets via scoped flooding.

Similarly to other reactive protocols, ODMRP consists of a request phase and a reply phase. When a multicast source has data to send but no route or group membership information is known, it piggybacks the data in a `Join-Query` packet. When a neighbor node receives a unique `Join-Query`, it records the upstream node ID in its *message cache*, which is used as the node's routing table, and rebroadcasts the packet. This process' side effect is to build the reverse path to the source. When a `Join-Query` packet reaches the multicast receiver, it generates a `Join-Table` packet that is broadcast to its neighbors. The `Join-Table` packet contains the multicast group address, sequence of (source address, next hop address) pairs, and a count of the number of pairs. When a node receives a `Join-Table`, it checks if the next node address of one of the entries matches its own address. If it does, the node realizes that it is on the path to the source and, thus, becomes a part of the forwarding group for that source by setting its *forwarding group flag*. It then broadcasts its own `Join-Table`, which contains matched entries. The next hop IP address can be obtained from the message cache. This process constructs (or updates) the routes from sources to receivers and builds the forwarding group. Membership and route information is updated by periodically (every `Join-Query-Refresh` interval) sending `Join-Query` packets. Nodes only forward (nonduplicate) data packets if they belong to the forwarding group or if they are multicast group members. By having forwarding group nodes flood data packets, ODMRP is more immune to link/node failures (e.g., due to node mobility). This is in fact an advantage of mesh-based protocols. Fig. 1 illustrates how the mesh is created in ODMRP.

2.2 Multicast Ad Hoc On-Demand Distance Vector (MAODV)

MAODV is an example of a tree-based multicast routing protocol (Fig. 2 illustrates MAODV tree formation). Similarly to ODMRP, MAODV creates routes on-demand. Route

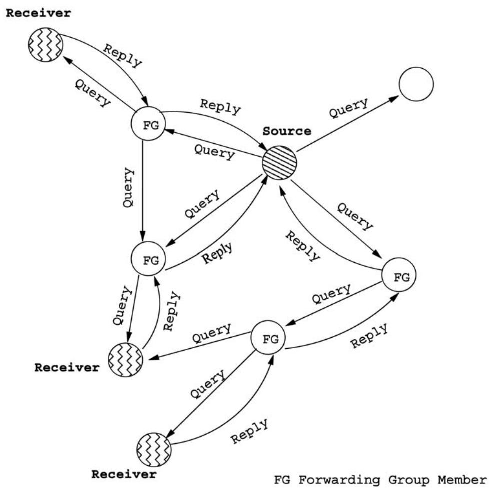

Fig. 1. Mesh formation in ODMRP.

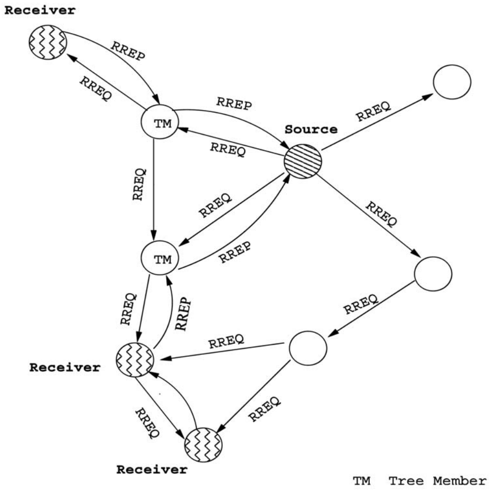

Fig. 2. Tree creation in MAODV.

discovery is based on a route request Rreq and route reply Rrep cycle. When a multicast source requires a route to a multicast group, it broadcasts a Rreq packet with the join flag set and the destination address set to the multicast group address. A member of the multicast tree with a current route to the destination responds to the request with a Rrep packet. Nonmembers rebroadcast the Rreq packet. Each node on receiving the Rreq updates its route table and records the sequence number and next hop information for the source node. This information is used to unicast the Rrep back to the source. If the source node receives multiple replies for its route request, it chooses the route having the freshest sequence number or the least hop count. It then sends a multicast activation message Mact which is used to activate the path from the source to the node sending the reply. If a source node does not receive a Mact message within a certain period, it broadcasts another Rreq. After a certain number of retries (Rreq-Retries), the source assumes that there are no other members of the tree that can be reached and declares itself the *Group Leader*. The group leader is responsible for periodically broadcasting group hello (Grp-Hello) messages to maintain group connectivity. Nodes also periodically broadcast Hello messages with *time-to-live = 1* to maintain local connectivity.

2.3 Flooding

Our implementation of routing by flooding is quite standard: When a node receives a packet, it broadcasts the packet except if it has seen that packet before. Nodes keep a cache of recently received packets; older packets are replaced by newly received ones. A node only rebroadcasts a packet if that packet is not in the node's cache.

We use a well-known randomization technique to avoid collisions: When a node receives a packet, it waits a random time interval between 0 and flooding interval before it rebroadcasts the packet.

Table 1 summarizes key characteristics of the three protocols under investigation.

3 Simulation Model and Methodology

We used ns-2 as the simulation platform. ns-2 is a popular discrete-event simulator which was originally designed for for wired networks and has been subsequently extended to support simulations in mobile wireless (and MANET) settings. In particular, we use the CMU Monarch group's extensions that enable ns-2 to simulate multihope MANETs [4]. Some MANET scenarios used in our simulations were generated using a scenario generator for ad hoc networks [18]; they are described in detail in Section 6.2 below.

3.1 MANET Scenarios

We use two type of MANET scenarios in our simulations. In "synthetic" scenarios, parameters such as mobility, multicast group size, traffic sources, and number of multicast groups are varied over an arbitrary range of values. We also define more "concrete" environments reflecting specific MANET applications, namely, impromptu teleconferencing and disaster relief/recovery scenarios. These more concrete MANET scenarios were generated using the scenario generator presented in [18] and are described in detail in Section 6.2.

In the synthetic scenario simulations, 50 nodes are randomly placed in a 1,000 m^2 field. Each node transmits a maximum of 1,000 packets (256 bytes each) at various times

TABLE 1
Protocol Summary

Protocol	Configuration	Loop Free	Periodic Messaging	Control Packet Flooding
Flooding	Mesh	Yes	No	No
ODMRP	Mesh	Yes	Yes	Yes
MAODV	Tree	Yes	Yes	Yes

TABLE 2
Simulation Parameters

Parameter	Value	Description
number-of-nodes	50	simulation nodes
num-packets	1000	messages sent by a node
packet-size	256 bytes	data packet size
field-range-x	1000 m	X-dimension of motion
field-range-y	1000 m	Y-dimension of motion
power-range	225 m	node's transmission range
bandwidth	2 Mbit/s	node's bandwidth
simulation-time	500 s	simulation duration
node-placement	random	node placement policy
propagation-func	Free-Space	propagation function
radio-type	Radio-No-Capture	no capture effect
mac-protocol	802.11	MAC layer
transport-protocol	UDP	transport layer

during the simulation. The nodes' channel bandwidth is set to 2 Mbit/sec and their transmission range is 225 meters. For these simulations, senders are chosen randomly from the multicast group members. All member nodes join at the start of the simulations and remain members throughout the duration of the simulation.

3.2 Mobility Model

The mobility model used is a modified version of the *random-waypoint* model we refer to as the *bouncing ball* model. In this model, nodes start off at random positions within the field. Each node then chooses a random direction and keeps moving in that direction until it hits the terrain boundary. Once the node reaches the boundary, it chooses another random direction and keeps moving in that direction until it hits the boundary again. An important aspect of our modified mobility model is that we always set *Vmin* to be nonzero. In fact, we set *Vmin* = *Vmax* for all synthetic scenarios. Hence, the bouncing ball model does not suffer from the drawbacks of the random mobility model as shown in [22].

3.3 Traffic Model

A constant bit rate (CBR) traffic generator was used for synthetic scenarios. The data payload size was fixed at 256 bytes. Senders were chosen randomly among network nodes. Network traffic for different sender populations was maintained constant at 50 Kbps by adjusting the inter-packet interval for the CBR sources. For concrete scenarios, we also used the ON-OFF traffic generator. Each source transmitted at 5 Kbps with a burst period of 3 secs and idle time of 3 secs.

3.4 Metrics

We use the following metrics in evaluating the performance of the different multicast routing protocols.

- **Packet delivery ratio** is computed as the ratio of total number of unique packets received by the receivers to the total number of packets transmitted by all sources times the number of receivers.
- **Routing overhead** is the ratio between the number of control bytes transmitted to the number of data bytes received. In ODMRP, control bytes account for `Join-Query` and `Join-Table` packets. It also includes data packet header bytes forwarded by forwarding group members. In MAODV, control bytes account for the `Rreq`, `Rrep`, `Mact`, `Hello`, and `Grp-Hello` packets. It also includes the data packet headers forwarded by intermediate nodes. In flooding, control bytes include all data header bytes forwarded by network nodes. We also account for the length of the IP header in our calculation of routing overhead.
- **Group reliability** is a measure of the effectiveness of the routing protocol in delivering packets to all receivers. We compute group reliability as the ratio of number of packets received by **all** multicast receivers to number of packets sent. Thus, for this metric, a packet is considered to be received only if it is received by every member of the multicast group.

Other Parameters. While Table 2 summarizes generic simulation parameters, Tables 3 and 4 summarize ODMRP and MAODV-specific parameters, respectively.

4 Simulation Results

In this section, we report simulation results comparing ODMRP, MAODV, and flooding. In these simulations, we use synthetic MANET scenarios, in which we subject the protocols to a wide range of mobility, traffic load, and multicast group characteristics (i.e., group size and number of sources). We ran each simulation (keeping all parameters constant) five times, each time using a different seed value. Each data point in the graphs below represents the average across all five runs. The error-bars shown in the graphs represent a confidence interval (CI) of 95 percent.[1]

In our simulations, the senders are chosen at random from the multicast group. All nodes join as members at the start of the simulations and remain members throughout the duration of the simulation.

We should point out that, in the performance study reported in [14], speeds were limited to 72 km/h and the

[1] Although we calculated the 95 percent CI for all graphs, we only show error-bars in graphs where they do not impact readability.

TABLE 3
ODMRP Parameters

Parameter	Value
Join Query refresh interval	3 secs
Forwarding Group Timeout	3 secs
Route Timeout	5 secs
Data Rebroadcast interval	25 ms

number of sources to 20 in a network of 50 nodes. Except when varying multicast group size from 5 to 40 receivers, all other simulations performed in the above-mentioned study used 20-node multicast groups. The study in [14] does not include MAODV in the pool of evaluated multicast protocols.

4.1 Effect of Mobility

The mobility experiment consisted of five traffic sources and 20 receivers chosen randomly. Each source transmitted 10 Kbps and, thus, the overall network load was 50 Kbps. Average node speed was varied between 0 and 150 kms/hr. Speeds of 150 kms/hr might at first seem too high. However, we claim that such high speed is very reasonable whenever a MANET includes fast-moving nodes, such as helicopters, fixed-wing aircraft, as well as police, military, and other emergency vehicles.

4.1.1 Packet Delivery Ratio

Fig. 3 shows how protocol reliability varies with mobility (node speed). The general trend we observe from the figure is that, especially at high mobility, flooding performs better than ODMRP, which in turn performs better than MAODV.

Comparing flooding to ODMRP, we notice that—at lower speeds—the difference in packet delivery ratio is between 5 percent and 7 percent. This result agrees with what was observed in [14]. However, at higher speeds, the gap in the packet delivery ratio starts widening. In the case of ODMRP, increased mobility requires that forwarding group members be updated more frequently. However, the frequency at which routes are refreshed (using periodic `Join-Queries`) remains constant, i.e., does not change with node speed. One way to address this problem is to update forwarding group members more often through more frequent *Join-Queries*. This, of course, would result in higher control overhead and, possibly, greater packet loss due to contention.

Comparing ODMRP with MAODV, we observe that ODMRP exhibits better (by roughly 10 percent) packet delivery ratios. Since ODMRP maintains meshes, it has multiple redundant paths to receivers and is not affected by mobility as greatly as MAODV. Increased mobility causes frequent link changes and requires MAODV to reconfigure the multicast tree more frequently to prevent stale routing information. This in turn requires higher control traffic, which can have a negative effect of increased packet loss due to contention and hidden terminals.

TABLE 4
MAODV Parameters

Parameter	Value
Group Hello Interval	5 secs
Hello Interval	1 sec
Mtree Build	3 secs
Route Discovery Timeout	3 secs

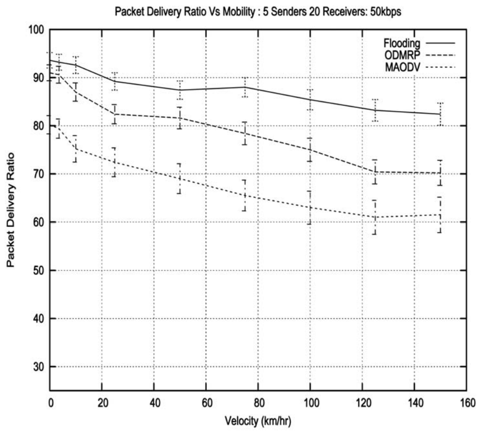

Fig. 3. Packet delivery ratio as a function of node mobility.

4.1.2 Routing Overhead

Fig. 4 plots control overhead per data byte transferred as a function of mobility. Note that flooding's overhead does not change with mobility as only data header packets contribute to overhead. In ODMRP, the *Join-Query* interval was fixed at 3 seconds and, hence, control overhead remains fairly constant with node mobility. The slight increase in overhead at higher speeds (around 55 km/hr) is due to the fact that the number of data bytes delivered decreases with increased mobility. In the case of MAODV, increased mobility causes frequent link breakages and data packet drops; link outages also generate repair messages, increasing control overhead.

4.1.3 Group Reliability

Since MANETs often target mission-critical applications, scenarios that require data transmission to be received by **all** multicast group members in a timely fashion are quite common. While a reliable transport protocol would repair losses detected by the communication end points, having the highest possible delivery rate from the underlying

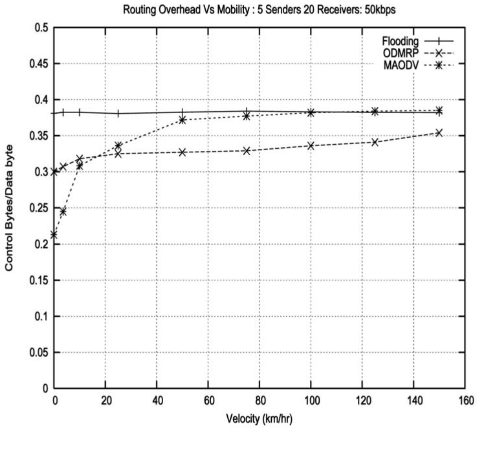

Fig. 4. Routing overhead as a function of node mobility.

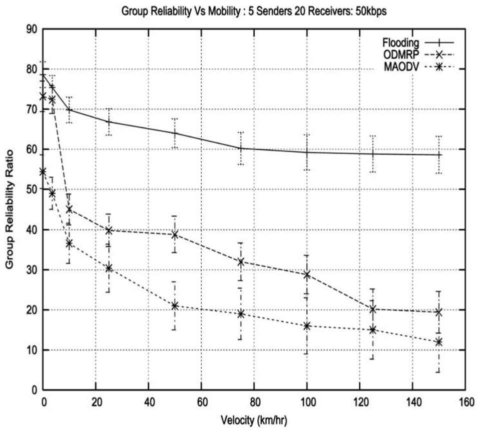

Fig. 5. Group reliability as a function of node mobility.

routing protocol improves the system's overall efficiency, including response time. Our group reliability metric tries to capture the effectiveness of routing protocols in delivering packets to all group members.

Fig. 5 plots group reliability as a function of node speed. From the figure, it can be seen that flooding is most effective in delivering packets to all group members (as expected). Moreover, flooding is able to keep group reliability fairly constant, even at higher speeds.

Both ODMRP and MAODV exhibit poor performance even at low mobility (group reliability lower than 50 percent for speeds higher than 10 km/hr). However, as expected, ODMRP exhibits better group reliability than MAODV. Although ODMRP can maintain multiple routes to receivers, the mesh connectivity is largely dependent on the number of senders and receivers. In the case of five senders, mesh connectivity is insufficient to ensure packet delivery to all group members (especially with node mobility), resulting in low group reliability.

Since MAODV delivers packets along a multicast tree, a single packet drop upstream can prevent a large number of downstream multicast receivers from receiving the packet. The absence of redundant routes affects performance greatly as node mobility results in frequent link breakages and packet drops.

4.2 Effect of Multicast Group Size

In this set of experiments, we focus on the influence of group size (the number of receivers) on multicast routing performance. The number of senders was fixed at 10, node mobility at 75 kms/hr, and traffic load at 50 Kbps. Group size was varied from 10-40 receivers in increments of 5.

4.2.1 *Packet Delivery Ratio*

Fig. 6 shows the variation in protocol reliability as a function of group size. Note that flooding is able to keep its delivery ratio fairly constant and close to 90 percent for different group sizes. Compared to ODMRP, flooding's delivery ratio is around 10 percent higher at a group size of 10 and around 6 percent higher as multicast group size increases to 40. Interestingly, ODMRP delivery ratio increases as group size increases. This is indeed consistent with the way mesh-based protocols operate. For instance, in ODMRP, the mesh is formed as a result of the `Join-Query-Join-Table` process. As the number of receivers increase, the number of `Join-Tables` sent out in response to `Join-Queries` increases. This causes a larger number of nodes to be incorporated into the mesh as *forwarding group* members, increasing mesh connectivity and redundancy. Hence, the packet delivery ratio tends to increase with increase in multicast receivers.

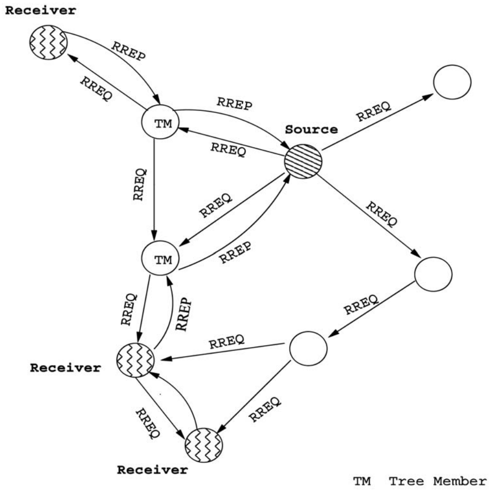

Fig. 6. Packet delivery ratio as a function of multicast group size.

In the case of MAODV, the packet delivery ratio decreases as group size increases (it is around 77 percent for 10 receivers and lowers to approximately 62 percent for 40 receivers). One reason for the decrease is that, as previously mentioned, a packet loss upstream affects a larger set of receivers. The increased group size also results in a greater number of control messages transmitted, which can result in greater packet loss due to collisions.

4.2.2 *Routing Overhead*

Fig. 7 shows how control overhead varies with group size. At low values of group size, flooding exhibits the highest routing overhead among all protocols for groups with up to 25 receivers. Flooding's overhead decreases with increasing group size. This is because all nodes rebroadcast data packets irrespective of group size. However, rebroadcast packets become more effective as group size increases since they now count toward packets delivered to multicast receivers. For this particular scenario, ODMRP's routing overhead is the highest among all three protocols for group sizes above 25. This is due to the large number of `Join-Tables` being transmitted and greater redundancy as the number of group members increases. In the case of MAODV, increased group size results in a larger number of `Repair` messages. However, data packets do not have to travel over multiple redundant paths, resulting in a lower overall routing overhead for MAODV as compared to ODMRP and flooding.

4.2.3 *Group Reliability*

Fig. 8 shows how group reliability varies with group size. As expected, group reliability of all protocol degrades for

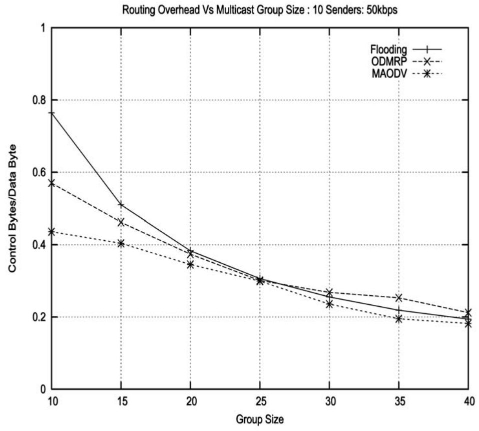

Fig. 7. Routing overhead as a function of group size.

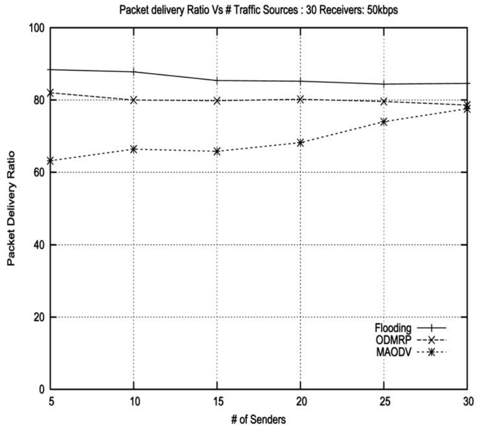

Fig. 9. Packet delivery ratio as a function of number of traffic sources.

larger multicast group size. This can be explained by the fact that, as the number of receivers increase, the probability of at least one receiver not receiving the data packet also increases.

From the graph, it is seen that the trend is similar to that observed in Section 4.1.3.

4.3 Effect of Number of Traffic Sources

In this set of experiments, we vary the number of multicast sources from 10 to 30 in steps of 5, keeping the number of receivers fixed at 30 and node mobility fixed at 75 kms/hr. For each value of number of senders, overall traffic load is maintained constant at 50 Kbps by changing the CBR sources' interpacket interval.

4.3.1 Packet Delivery Ratio

Fig. 9 shows the packet delivery ratio as a function of the number of senders. Note that both the flooding and ODMRP packet delivery ratios remain fairly constant with the number of senders; thus, they do not suffer from increased contention except at a higher number of sources, where a slight drop off can be observed and is attributed to data packet loss due to collisions. An interesting and counter intuitive result is that, in the case of MAODV, the delivery ratio increases with an increase in the number of traffic sources. This is due to the fact that, in MAODV, the shared tree is formed as a result of the `Rreq-Rrep` process. As the number of senders increases, a greater number of intermediate nodes (on the path from the sender to the multicast tree) are grafted as part of the tree. This helps to increase redundancy along certain links due to the presence of multiple downstream neighbors who can potentially forward data along the tree. Hence, the packet delivery ratio tends to increase with an increase in the number of sources. However, MAODV packet delivery ratio is consistently lower than that of ODMRP and flooding.

4.3.2 Routing Overhead

Fig. 10 depicts how control overhead varies with the number of traffic sources. Flooding does not transmit any control messages and, hence, its routing overhead remains constant with the number of senders. For ODMRP, increased sender population results in a larger number of `Join-Reqs` and `Join-Tables`. `Join-Tables` in particular can result in large byte overhead since they carry next-hop information for multiple sources. Similarly, larger sender population results in a larger number of MAODV control messages being transmitted. However, as discussed in Section 4.3.1, the number of data bytes received also increases. Hence, MAODV's overall ratio of control bytes/data byte delivered remains fairly constant.

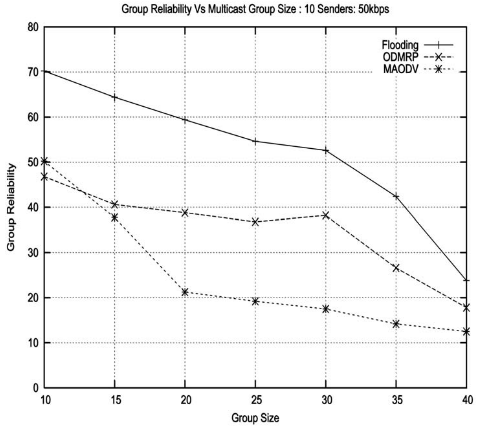

Fig. 8. Group reliability as a function of group size.

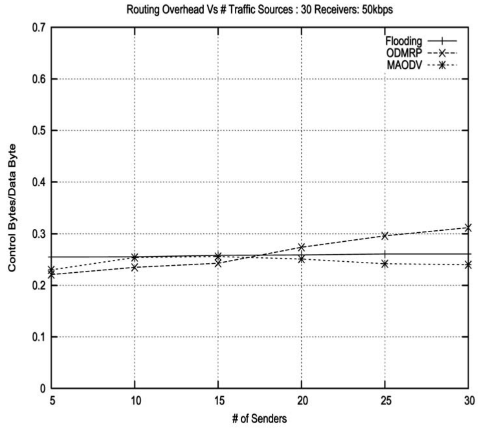

Fig. 10. Routing overhead as a function of traffic sources.

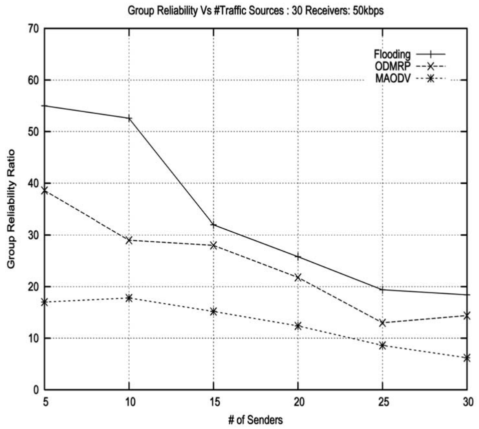

Fig. 11. Group reliability as a function of traffic sources.

4.3.3 Group Reliability

Fig. 11 shows how group reliability varies with the number of traffic sources. It is interesting to notice how the different reliability metrics capture different protocol behavior. According to the packet delivery ratio metric, both flooding and MAODV exhibit fairly high delivery ratios (above 80 percent); MAODV delivers around 65 percent of the packets for up to 20 senders, but increases its reliability (close to 80 percent) for 30 receivers, as shown in Fig. 9. However, the group reliability metric, as depicted in Fig. 11, shows a completely different behavior. Even though the relative performance among the protocols remains the same (i.e., $flooding >> ODMRP >> MAODV$, where $>>$ denotes "performs better"), we observe that group reliability degrades considerably for larger number of senders. This effect is mainly due to increased contention as a larger number of senders results in a greater number of packets transmitted. As a result, a greater number of packets are dropped due to collisions.

4.4 Multiple Multicast Groups

The goal of these experiments is to evaluate how multiple multicast groups impact the performance of mesh and tree-based multicast routing. For the multigroup simulations, two separate multicast groups are used, each of which has five sources and 10 receivers. Average node speed and overall traffic load are fixed at 20 Km/hr and 50 Kbps, respectively. For the single-group simulations, we use 10 senders and 20 receivers. The same node mobility and overall traffic load are used, i.e., 20 Km/hr and 50 Kbps, respectively.

Table 5 compares the performance of the protocols when operating in a multigroup environment against single multicast group operation. We observe that flooding's performance is the most affected by multiple multicast groups. Although the delivery ratio remains fairly similar, routing overhead almost doubles. This is due to the fact that, since flooding does not maintain group membership information, nodes rebroadcast every packet irrespective of the group.

In the case of ODMRP, mesh connectivity depends on the number of receivers. Since, in the multiple group case, the number of receivers for each group is halved (as compared to the single group case), the mesh is not as rich as before, resulting in lower packet delivery ratios. Routing overhead decreases since nodes can piggyback the `Join-Tables` for multiple groups. The performance of MAODV is not significantly affected by multigroup operation.

4.5 Effect of Network Traffic Load

In this section, we evaluate the impact of increasing traffic load on protocol performance. The number of senders was fixed at 10 and number of receivers at 20, respectively. Node mobility was set at 75 kms/hr. The overall network load was increased from 10 Kbps to 50 Kbps in steps of 5 Kbps. This is achieved by increasing the sending rate of each source from 1 Kbps to 5 Kbps. The data traffic introduced into the network is CBR traffic.

4.5.1 Packet Delivery Ratio

Fig. 12 shows the packet delivery ratio as a function of traffic load. It is observed that all protocols are affected by the increase in network traffic. Increased network traffic results in greater contention and packet loss due to higher collisions and buffer overflow. For the traffic loads considered, flooding still outperforms ODMRP and MAODV in terms of delivery ratios. However, we expect the performance of flooding to deteriorate more rapidly than ODMRP and MAODV as traffic load increases on account of the greater number of redundant transmissions.

4.5.2 Routing Overhead

Fig. 13 depicts the control overhead per data byte delivered as a function of traffic load. It can be seen that flooding's control overhead remains almost constant with increasing load. Flooding does not transmit any control packets and all packets received by a node are retransmitted exactly once, resulting in almost constant control overhead. The high routing overhead seems to suggest that flooding can be quite expensive at higher traffic loads and, hence, not scalable with increased traffic loads.

TABLE 5
Performance with Multiple Multicast Groups

Protocol	Pkt Delivery Ratio	Routing Overhead	Group Reliability
Flooding (1 Group)	87.6	0.383	59.42
Flooding (2 Groups)	86.8	0.764	70.41
ODMRP (1 Group)	79.6	0.374	38.80
ODMRP (2 Groups)	71.8	0.328	36.27
MAODV (1 Group)	70.2	0.345	21.25
MAODV (2 Groups)	68.2	0.352	23.52

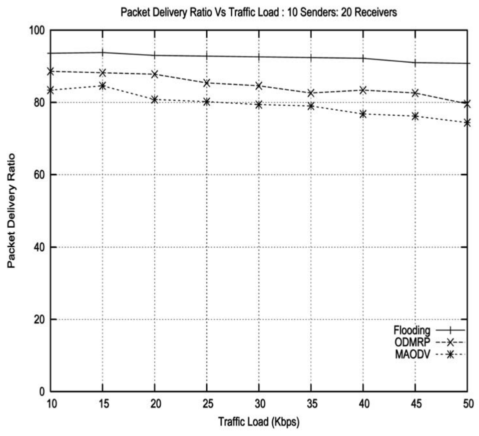

Fig. 12. Packet delivery ratio as a function of traffic load.

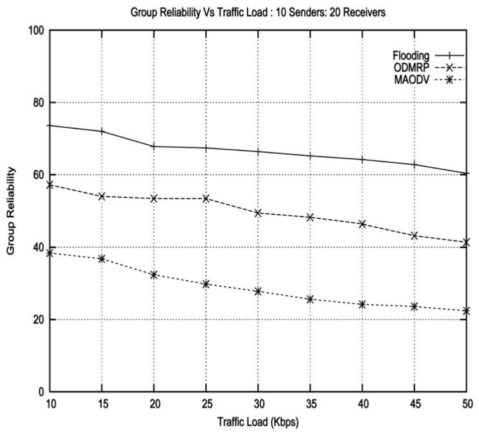

Fig. 14. Group reliability as a function of traffic load.

In the case of ODMRP and MAODV, routing overhead decreases with an increase in traffic load. As network load increases, the total number of data bytes received by ODMRP and MAODV receivers also increases. However, control data transmitted remains fairly constant with increased network load, thereby reducing the routing overhead. (Note that routing overhead is calculated as the ratio of control bytes/data byte received). In this experiment, ODMRP has a greater routing overhead than MAODV because of the mesh structure, but the gap reduces as the network load increases. As traffic load increases, both ODMRP and MAODV are affected by packet losses because of contention. Since ODMRP maintains multiple routes to destinations, receivers can possibly receive data packets from other routes. This increases the total number of data bytes received by ODMRP receivers as compared to MAODV receivers, which helps to reduce the routing overhead.

4.5.3 Group Reliability

Fig. 14 plots group reliability as a function of traffic load. From the figure, it can be seen that group reliability for all protocols decreases with increase in traffic load as expected. Flooding has the highest group reliability among the three protocols, as before. All protocols exhibit an almost similar decrease in group reliability (about 16-18 percent) as traffic load increases. In the case of flooding and ODMRP, the increased redundancy is offset by the increase in collisions, which degrades the reliability. In the case of MAODV, performance degrades because of increased collisions and buffer overflow as traffic load increases.

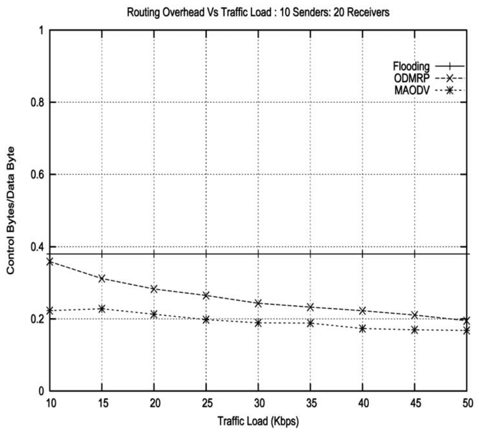

Fig. 13. Routing overhead as a function of traffic load.

4.6 Qualitative Comparison of Protocols

Table 6 provides a qualitative comparison of the protocols based on our simulation analysis in the preceding sections.

Flooding requires no resources for route initialization since there is no setup associated with establishing routes to multicast group members. In the case of ODMRP, nodes have to transmit `Join-Query` messages to establish routes to multicast group members. Group members and forwarding group nodes reply with `Join-Tables`. MAODV's route initialization consists of leader selection for the group followed by a `Rreq-Rrep` route discovery phase. The sender then has to send out a `Mact` message to activate a particular route among various possible routes. Thus, MAODV incurs the highest overhead for route setup.

In terms of data forwarding, as observed from the simulation results, flooding has the highest overhead for most scenarios. Hence, bandwidth resources used by flooding in delivering data to receivers is greatest among the protocols considered. MAODV has the lowest forwarding overhead, whereas forwarding resources used by ODMRP are moderate.

As expected, flooding delivers the highest reliability (both in terms of packet delivery ratio and group reliability). Since ODMRP maintains a mesh structure and has multiple routes to multicast group members, it exhibits better reliability than MAODV, but lower than that of flooding. MAODV, on the other hand, maintains a shared tree structure and is susceptible to frequent link changes due to mobility. This has a considerable effect on MAODV's reliability.

In flooding, data is rebroadcast by all nodes and does not travel along certain paths, resulting in low traffic concentration on any given link. Because of the mesh structure in ODMRP, data is routed through multiple paths. In the case of MAODV, data has to be forwarded along the tree and can lead to traffic concentration along certain tree links.

Flooding is not very scalable with an increase in the number of nodes because of the excessive broadcasts and forwarding overhead. In the case of ODMRP, routing

TABLE 6
Qualitative Comparison of ODMRP, MAODV, and Flooding

Protocol	Route setup overhead	Route maintenance overhead	Data forwarding overhead	Reliability	Traffic concentration	Scalability
Flooding	Low	Low	High	High	Low	Low
ODMRP	Moderate	Moderate	Moderate	High	Low	Low
MAODV	High	Highest	Low	Low	Highest	High

overhead can get prohibitive as the number of sender increases. MAODV is most scalable in terms of the number of network nodes and multicast senders.

5 Flooding Variations

Our simulation results show that flooding has higher reliability compared to ODMRP and MAODV, especially at high mobility and traffic load. However, one major drawback of flooding is redundant broadcasts, which can considerably increase data forwarding overhead. Redundant broadcasts are particularly damaging in ad hoc networks where nodes are often bandwidth and energy-constrained. In this section, we introduce *scoped flooding*, a variation of flooding that aims at restricting redundant broadcasts. It does so based on different heuristics, which are discussed in detail below.

It is also possible to envisage scenarios that require higher delivery guarantees beyond what plain flooding can provide. To achieve these more stringent delivery guarantees, we propose a technique called *hyper flooding*. The basic principle of hyper flooding is to force nodes to rebroadcast data packets more than once based on certain criteria. This helps to ensure maximum packet delivery at the cost of overhead. We argue that mission critical applications may be willing to pay the price of higher overhead in exchange for the highest possible delivery guarantees. Below, we describe scoped and hyper flooding in detail.

5.1 Scoped Flooding

The basic principle behind scoped flooding is the reduction of rebroadcasts to avoid collisions and minimize overhead. Scoped flooding is suitable for constrained mobility environments (e.g., conference scenarios) where nodes do not move much and, thus, plain flooding will likely yield unnecessary redundant rebroadcasts. In fact, Ni et al. [15] show that the coverage area of subsequent retransmissions reduces drastically and drops down to 0.05 percent when the number of retransmissions is greater than four.

Different heuristics can be used in deciding whether to rebroadcast a packet. In our scoped flooding implementation, each node periodically transmits `hello` messages which also contain the node's neighbor list. Nodes use `hello` messages to update their own neighbor list and add received lists to their neighbor list table. When a node receives a broadcast, it compares the neighbor list of the transmitting node to its own neighbor list. If the receiving node's neighbor list is a subset of the transmitting node's neighbor list, then it does not rebroadcast the packet. In our simulations, we did not require neighbor lists to be strict subsets of one another. An 85 percent overlap was considered sufficient to prevent Rebroadcasts.[2]

5.2 Hyper Flooding

Hyper flooding is suitable for highly mobile scenarios where high reliability is required. The price to pay for the additional reliability is, of course, higher overhead.

Similarly to both plain and scoped flooding, nodes in hyper flooding exchange periodic `hello` messages. When a node receives a `hello` message from a neighbor, it adds the identity of the `hello` message originator to its neighbor list (if the list does not already contain that node). As in plain flooding, when a node receives a new data packet, it simply rebroadcasts the packet and queues it in its packet cache. Additionally, rebroadcasts are triggered by two other events: receiving a data packet from a node which is not in the current neighbor list or receiving a `hello` message from a new neighbor. In these cases, nodes transmit all packets in their cache. The rationale behind rebroadcasts is that "newly acquired" neighbors could have missed the original flooding wave because of their mobility. This increases overall reliability by ensuring that new nodes entering the transmission region of a node receive data packets which they otherwise would have missed. Nodes periodically purge their packet cache to prevent excess reflooding of older packets.

6 Performance of Flooding Variations

We conducted extensive simulations to compare the performance of the proposed flooding variations against plain flooding, ODMRP, and MAODV. One novel feature of our study is that, in addition to the synthetic environments described in Section 3, we also use concrete MANET scenarios, namely, conferencing and emergency response/rescue operations (described in detail in Section 6.2 below). We start with the simulation results for synthetic MANET scenarios.

6.1 Synthetic Scenarios

Similarly to the scenarios described in Section 3, for these simulations, 150 nodes are randomly placed in a 1,500 m^2 field. Each node transmits a maximum of 1,000 packets (256 bytes each) at various times during the simulation. The nodes' channel bandwidth is set to 2 Mbit/sec and their transmission range is 225 meters. Senders are chosen randomly from the multicast group members. All member nodes join at the start of the simulations and remain members throughout the duration of the simulation. Total

[2] This value was chosen after extensive simulation-based analysis of scoped flooding.

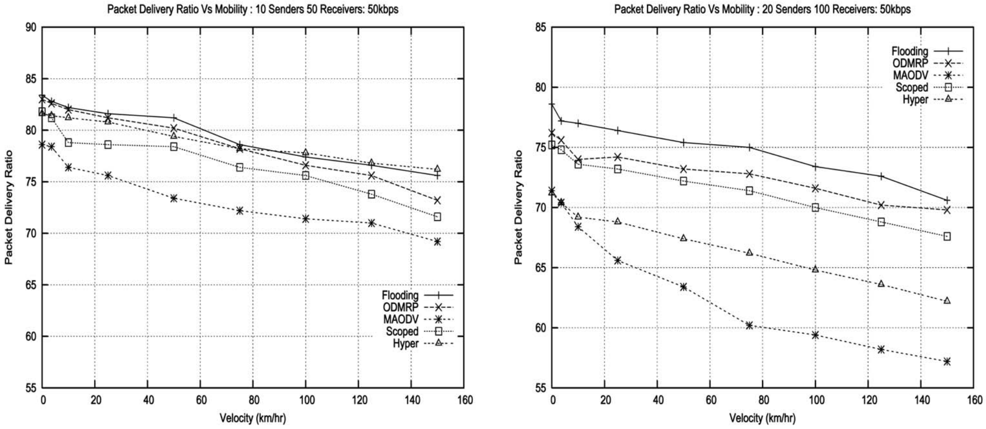

(a) (b)

Fig. 15. Packet delivery ratio as a function of node mobility. (a) Ten senders, 50 receivers. (b) Twenty senders, 100 receivers.

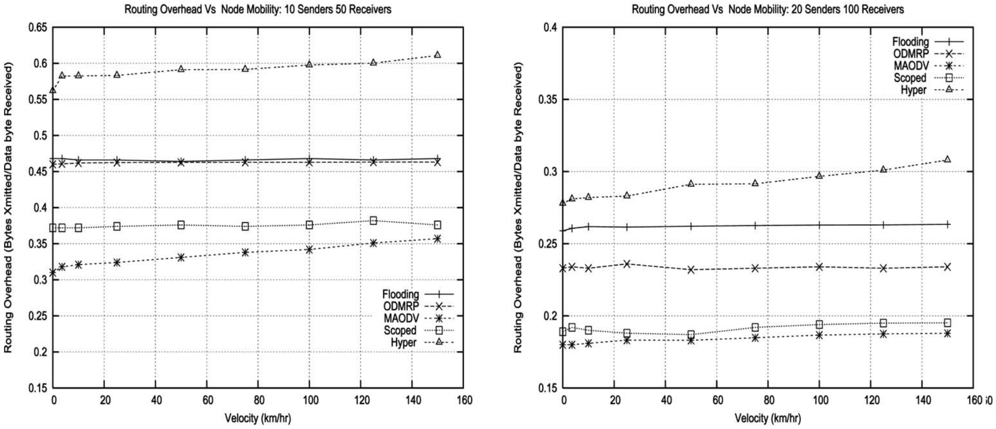

(a) (b)

Fig. 16. Control overhead as a function of node mobility. (a) Ten senders, 50 receivers. (b) Twenty senders, 100 receivers.

network traffic was kept constant at 50 Kbps. Each data point was obtained by averaging across five runs with different seed values. Although we ran simulations for 50 nodes in a 1,000 m^2 field, the results obtained are very similar to those presented in Section 4.1. Hence, we do not present those results again.

Fig. 15 shows how packet delivery ratio varies with mobility and the number of traffic sources. Surprisingly, for these scenarios, hyper flooding does not exhibit the highest delivery ratio among all protocols as expected. Given the larger node density for these particular scenarios, rebroadcasting multiple times seems counter effective, resulting in a lot of packets being dropped due to collisions. It is seen that flooding has the best delivery ratio, outperforming ODMRP by 2-3 percent.

However, this increase in reliability is obtained at the cost of routing overhead, as evident from Fig. 16. Another interesting observation is that the delivery ratio of scoped flooding is very similar to ODMRP. However, this reliability is obtained at a much lower routing overhead. Both ODMRP and scoped flooding have multiple redundant routes to destinations. However, in the case of scoped flooding, the number of redundant broadcasts is optimized by using forwarding nodes with nonoverlapping neighbors. Another factor contributing to scoped flooding outperforming ODMRP is that scoped flooding does not have to transmit any control messages which can potentially result in medium contention and higher packet loss due to collisions.

Fig. 17 plots group reliability as a function of node speed. From the figure, it can be seen that the protocols perform quite poorly in terms of delivering packets to all group members, especially at high mobility. The group reliability for all protocols drops to about 10-15 percent at speeds of 150 kms/hr with MAODV having the lowest group reliability. For these scenarios, ODMRP has the highest group reliability among the protocols evaluated. In these experiments, given the large node density and receiver population, flooding and hyper are severely affected by packet losses due to collision and contention. ODMRP and scoped flooding perform the best under these conditions because of their limited rebroadcasts as compared to flooding and hyper flooding. Since MAODV maintains a shared tree structure, it is susceptible to frequent link breakages due to mobility. This has a severe effect on MAODV's group reliability performance.

6.2 Concrete Scenarios

We also use "typical" MANET scenarios such as conferencing and rescue operations to compare the performance of

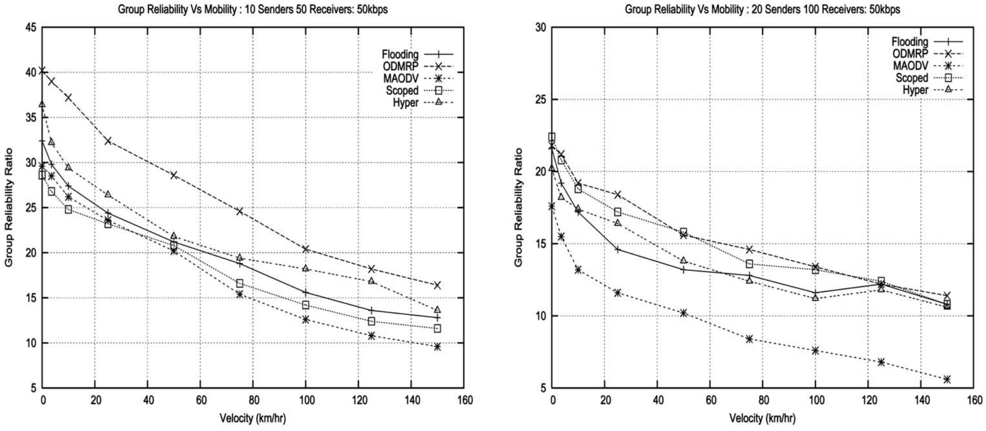

(a)

(b)

Fig. 17. Group reliability as a function of node mobility. (a) Ten senders, 50 receivers. (b) Twenty senders, 100 receivers.

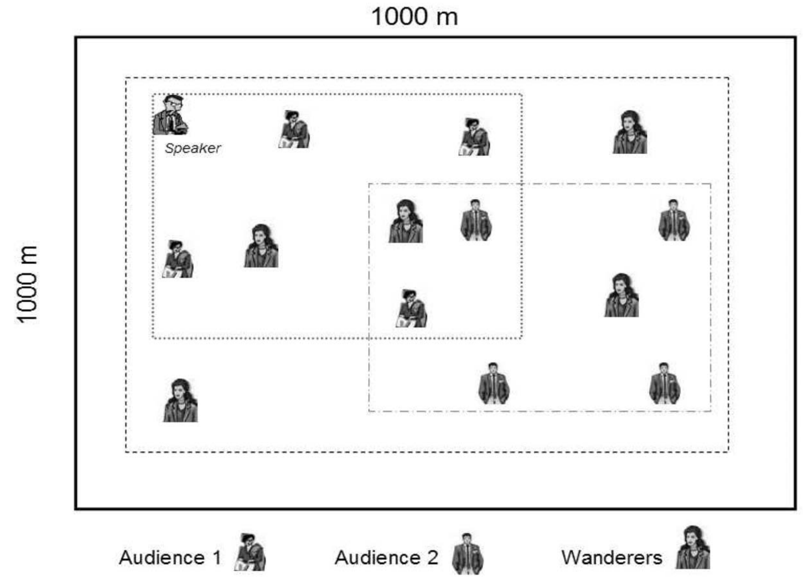

Fig. 18. Conference scenario setup.

the protocols under investigation. Such scenarios were generated using the scenario generator for ad hoc networks [18] and are described in greater detail below.

6.2.1 Conferencing

The conference scenario consists of a total of 50 nodes in a 1,000 m^2 field with one *speaker* node and three *audience* groups, i.e., *audience1*, *audience2*, and the *wanderers*. Both *audience1* and *audience2* consist of 20 members moving at low speeds (between 2-5 m/s) with pause time between 0-2 secs. The movement of the speaker was modeled using brownian motion, whereas the movement of the audience groups was modeled using random waypoint motion and node movement was restricted to a limited area within the field. *Wanderers* consist of nine nodes that were capable of moving over the entire topology. The speeds for these nodes were randomly chosen between 1-5 m/s with pause times between 0-1 sec. *Wanderers* move according to the random waypoint model. The *speaker* node and 20 randomly chosen audience nodes acted as sources of data.

Both CBR and ON-OFF traffic were used. In CBR, each source transmitted 2.5 Kbs, while the traffic rate was set to 5 Kbs for ON-OFF traffic with a burst period of 3 secs and idle time of 3 secs. Fig. 18 depicts the conference scenario setup.

Table 7 summarizes simulation results for the conferencing scenario in decreasing order of packet delivery ratio. Scoped flooding is the best performer for both CBR and ON-OFF traffic. In particular, for ON-OFF traffic, scoped flooding's delivery ratio is around 10 percent higher than ODMRP and around 14 percent higher than MAODV, yet its overhead is lower than ODMRP and only slightly higher than MAODV. Flooding and hyper flooding exhibit lower delivery ratio than ODMRP and MAODV for CBR traffic. The low mobility of nodes coupled with sufficiently high node density and high traffic load results in a large number of collisions, especially for flooding and hyper flooding. The high overhead incurred by both protocols also contributes to increased medium contention.

TABLE 7
Conference Scenario

Conference scenario			
	Protocol	Delivery ratio %	Routing overhead (Bytes Xmitted/data bytes recvd)
CBR Traffic	Scoped flooding	84.2	0.114
	ODMRP	81.4	0.136
	MAODV	76.8	0.081
	Hyper flooding	71.2	0.145
	Flooding	70.6	0.137
ON-OFF Traffic	Scoped flooding	76.4	0.126
	ODMRP	67.3	0.128
	Flooding	64.5	0.154
	MAODV	63.5	0.084
	Hyper flooding	60.4	0.172

TABLE 8
Emergency Response Scenario

Emergency response scenario			
	Protocol	Delivery ratio %	Routing overhead (Bytes Xmitted/data bytes recvd)
CBR Traffic	Hyper flooding	80.2	0.148
	Flooding	76.4	0.132
	Scoped flooding	75.2	0.116
	ODMRP	67.4	0.126
	MAODV	60.2	0.091
ON-OFF Traffic	Hyper flooding	78.4	0.165
	Flooding	73.2	0.141
	Scoped flooding	69.8	0.122
	ODMRP	60.36	0.129
	MAODV	56.2	0.093

6.2.2 Emergency Response Scenario

For the emergency response scenario, we use a 2,000 m^2 field with a total of 75 nodes divided into the following categories: two helicopters, two rescue teams of ground personnel, and two teams on ground vehicles. The helicopters move with speeds ranging between 0-50 m/s according to the random waypoint model. The first vehicle team consists of 25 nodes, while the second team consisted of eight nodes. Members of both vehicle teams move according to the random waypoint model with speeds ranging between 5-15 m/sec. The team of ground personnel consists of 20 nodes moving with speeds ranging between 0-5 m/s and pause times between 0-2 secs. Each team covers well-defined areas within the field with sufficient overlap to ensure that information could be relayed among the different teams. Two helicopters and 20 other randomly chosen nodes act as data sources for this scenario.

From Table 8, which summarizes simulation results for the emergency response scenario, we observe that flooding variations achieved considerably better packet delivery ratio than ODMRP or MAODV for both CBR and ON-OFF traffic. Even though we ensure that different “mobility” groups have sufficient overlap to relay data among groups, in the case of ODMRP, only forwarding group members can relay data, whereas, in MAODV, only multicast tree members can forward data traffic. At route setup time, nodes in the overlap region are incorporated as forwarding group members (ODMRP) or multicast tree members (MAODV). However, node mobility may cause forwarding group members and multicast tree members to move outside the overlap region, resulting in a large number of packet drops until the route is refreshed at the end of the *Active-Route-Interval*. This effect is more severe for bursty traffic as compared to CBR traffic. In the case of flooding and its variations, all nodes can forward data traffic and, thus, achieve better reliability. In particular, scoped flooding achieves close to 10 percent higher delivery ratio than ODMRP at lower overhead; when compared to MAODV, scoped flooding delivers close to 15 percent more packets at slightly higher overhead. Hyper flooding improves reliable delivery even further: For both CBR and ON-OFF traffic, it achieves between 20-22 percent better reliability than MAODV incurring approximately 50 percent overhead increase. When compared to ODMRP, hyper flooding’s reliability improvement is also quite substantial at slightly higher routing overhead.

7 Related Work

Since group-oriented services have been recognized as one of the primary applications of MANETs, several MANET multicast routing protocols have been proposed. As previously discussed, MANET multicast protocols can be classified according to how they propagate data as tree-based or mesh-based. While tree-based protocols propagate data over a tree spanning all multicast group members, in mesh-based protocols, a subset of network nodes (the mesh) is responsible for forwarding data to all multicast receivers. MANET protocols can also be classified according to how they acquire/maintain routes. Reactive (or on-demand) protocols acquire routes on demand and proactive protocols maintain routing state. Examples of mesh-based protocols include ODMRP, CAMP [6], and flooding. ODMRP is reactive, while CAMP is proactive. Flooding, of course, does not require routing state maintenance. AMRoute [2] and AMRIS [21] are examples of proactive, tree-based protocols. MAODV exemplifies a reactive, tree-based protocol. In addition to the above schemes, we highlight the Zone Routing Protocol (ZRP) [17] which uses a hybrid approach combining proactive route maintenance among nodes within a zone and reactive routing for interzone communication. Hybrid routing schemes are particularly useful for scenarios in which ad hoc networks may be connected to the wired infrastructure through gateway nodes. In such cases, member nodes requiring connectivity to the wired network may have to maintain routes to the gateway nodes at all times.

With the advent of GPS (Global Positioning Systems), protocols which make use of location information have been proposed. Knowledge of node positioning can help to make routing more effective at the cost of updating location information. Examples of such protocols are Location Aided Routing (LAR) [11], Zone-Based Hierarchical Link State (ZHLS) [9], and Distance Routing Effect Algorithm for Mobility (DREAM) [1].

Several studies have evaluated the performance of **unicast** routing protocols for MANETs [3], [5], [13], [10]. The only performance study of MANET **multicast** routing protocols that preceded our work was reported in [14]. This comparative study analyzed five different protocols, including ODMRP and flooding. It concluded that mesh-based protocols in general and ODMRP in particular perform better than tree-based approaches. As we previously pointed out, this work focused on a different portion of the multicast routing protocol design space. It evaluated protocols for lower mobility and traffic load scenarios and involving smaller sets of traffic sources. For these scenarios, ODMRP outperformed all the other protocols evaluated, except for flooding. As we confirmed in our study, for these scenarios, the performance difference between flooding and ODMRP is reasonably small.

More recently, Kunz and Cheng [12] published a comparative study evaluating the performance of ODMRP and MAODV under various scenarios. However, the results from their paper cannot be validated since they use the random waypoint model with Vmin = 0 m/s. It has already been shown in [22] that the random waypoint model with Vmin = 0 m/s can cause the average network mobility to tend to zero at steady state, leading to inaccurate simulation results.

Our study targets a different segment of the design space. It subjects mesh and tree-based protocols to a wider range of multicast group characteristics and network conditions, including more stringent mobility and traffic load scenarios. Besides a quantitative analysis, we also provide a qualitative comparison between mesh and tree-based multicast for MANETs. Based on our comparative analysis, we introduce two flooding variations, one which aims at reducing flooding's overhead and the other at improving flooding's reliability. Another contribution of our study is the use of concrete environments in evaluating MANET protocols.

8 Conclusions

In this paper, we reported on simulation-based experiments evaluating two different approaches to multicast communication in mobile ad hoc networks (MANETs), namely, mesh and tree-based multicast. One of the chief contributions of this work is our objective analysis of these two multicast routing protocol categories in order to characterize their behavior under a wide range of MANET scenarios, including different mobility and traffic load conditions as well as multicast group characteristics (e.g., size, number of sources, multiple multicast groups, etc.). Another contribution of this paper is the use of realistic MANET scenarios, such as conferencing and emergency response, in evaluating routing protocols. These MANET scenarios were generated using the scenario generator tool [18].

Our simulation results demonstrate that, even though the performance of all multicast protocols degrade, in terms of packet delivery and group reliability as node mobility and traffic load increases, mesh-based protocols (e.g., flooding and ODMRP) perform considerably better than tree-based protocols (e.g., MAODV). The general conclusion from the comparative analysis was that flooding, which is the simplest routing mechanism, provides higher delivery guarantees than ODMRP and MAODV for most scenarios considered. ODMRP exhibits decent robustness because of its mesh structure. MAODV did not perform as well as the other protocols in terms of packet delivery ratio and group reliability, but has the lowest routing overhead among the protocols considered.

A well-known drawback of flooding is its inherent overhead in the form of redundant broadcasts. This is particularly evident in the case of multiple multicast groups, where flooding's overhead increases with the number of groups. To limit flooding's excessive overhead, we proposed *scoped flooding*, a variation of flooding which attempts to minimize rebroadcasts by using neighbor information. Simulation results show that scoped flooding can reduce overhead by around 20 percent compared to flooding and 15 percent compared to ODMRP at comparable delivery ratios. One interesting observation was the performance of scoped flooding in conference scenarios, where it exhibited stellar performance in delivering data at low routing overhead.

In order to address the issue of reliability at high node speeds, we also investigated an other flooding variation, referred to as *hyper flooding*. Simulations results indicate that hyper flooding, indeed provides the best delivery guarantees under more stringent conditions (e.g., high mobility, traffic load), but this is achieved at greater overhead (about 10 percent in the case of our emergency response scenarios) than flooding. However, we believe that hyper flooding can be justified in those MANET scenarios demanding the highest possible guarantees of reliable (yet timely) delivery, regardless of costs.

One of the conclusions from our study is that, given the diversity of MANETs, it is impossible for any one routing protocol to be optimal under all scenarios and operating conditions. One possible solution would be to develop specialized multicast solutions for each type of network and the means for integrating those solutions. We believe that an adaptive, integrated approach to routing may be the best means to tackle this problem. In this approach, nodes can dynamically switch routing mechanisms based on their perception of the network conditions. However, such an adaptive approach presents various challenges such as: 1) Interoperability and integration issues and 2) mechanisms for active, on-the-fly switching among different multicast routing mechanisms as a mobile host changes the network type it is part of.

We are currently investigating adaptive, integrated approaches to MANET routing. At the same time, we are also developing analytical models for flooding mechanisms in order to obtain performance bounds and use those models as means to validate our simulation results.

REFERENCES

[1] S. Basagni, I. Chlamtac, V.R. Syrotiuk, and B.A. Woodward, "A Distance Routing Effect Algorithm for Mobility (Dream)," *Proc. IEEE/ACM MOBICOM '98 Conf.,* pp. 76-84, Oct. 1998.

[2] E. Bommaiah, M. Liu, A. McAuley, and R. Talpade, "AMRoute: Adhoc Multicast Routing Protocol," IETF manet, draft-talpade-manet-amroute-00.txt, Aug. 1998.

[3] J. Broch, D.A. Maltz, D.B. Johnson, Y.C Hu, and J. Jetcheva, "A Performance Comparison of Multihop Wireless Ad Hoc Network Routing Protocols," *Proc. ACM/IEEE MOBICOM '98 Conf.,* pp. 85-97, 1998.

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

[5] S.R. Das, R. Castaneda, J. Yan, and R. Sengupta, "Comparative Performance Evaluation of Routing Protocols for Mobile, Ad Hoc Networks," *Proc. IEEE Seventh Int'l Conf. Computer Comm. and Networks,* pp. 153-161, 1998.

[6] J. Garcia-Luna-Aceves and E. Madruga, "A Multicast Routing Protocol for Ad Hoc Networks," *Proc. INFOCOM '99 Conf.,* pp. 784-792, Mar. 1999.

[7] M. Gerla and S.J. Lee, "On-Demand Multicast Routing Protocol for Mobile Ad Hoc Networks," http://www.cs.ucla.edu/NRL/wireless/, 2005.

[8] P. Jacquet, P. Minet, A. Laouiti, L. Viennot, T. Clausen, and C. Adjih, "Multicast Optimized Link State Routing," IETF manet, draft-ietf-manet-olsr-molsr-01.txt, 2002.

[9] M. Joa-ng and L.T. Lu, "A Peer-to-Peer Zone Based Two-Level Link State Routing for Mobile Ad Hoc Networks," *IEEE J. Selected Areas in Comm.,* vol. 17, no. 8, pp. 1415-1425, Aug. 1999.

[10] P. Johansson, T. Larsson, N. Hedman, B. Mielczarek, and M. Degermark, "Scenario-Based Performance Analysis of Routing Protocols for Mobile Ad Hoc Networks," *Proc. IEEE/ACM MOBICOM '99 Conf.,* pp. 195-206, 1999.

[11] Y.B. Ko and N.H. Vaidya, "Location-Aided Routing in Mobile Ad Hoc Networks," *Proc. IEEE/ACM MOBICOM '98 Conf.,* pp. 66-75, 1998.

[12] T. Kunz and E. Cheng, "Multicasting in Ad Hoc Networks: Comparing MAODV and ODMRP," *Proc. Workshop Ad Hoc Comm.,* Sept. 2001.

[13] S.J. Lee, M. Gerla, and C.K. Toh, "A Simulation Study of Table-Driven and On-Demand Routing Protocols for Mobile Ad Hoc Networks," *IEEE Network,* vol. 13, no. 4, pp. 48-54, 1999.

[14] S.J. Lee, W. Su, J. Hsu, M. Gerla, and R. Bagrodia, "A Performance Comparison Study of Ad Hoc Wireless Multicast Protocols," *Proc. IEEE Infocom 2000 Conf.,* Mar. 2000.

[15] S. Ni, Y. Tseng, Y. Chen, and J. Sheu, "The Broadcast Storm Problem in a Mobile Ad Hoc Network," *Proc. IEEE/ACM MOBICOM '99 Conf.,* pp. 151-162, 1999.

[16] K. Obraczka, G. Tsudik, and K. Viswanath, "Pushing the Limits of Multicast in Ad Hoc Networks," *Proc. 21st Int'l Conf. Distributed Computing Systems,* pp. 719-722, Apr. 2001.

[17] M.R. Pearlman and Z.J. Haas, "Determining the Optimal Config. Ation for the Zone Routing Protocol," *IEEE J. Selected Areas in Comm.,* vol. 17, no. 8, pp. 1395-1414, Aug. 1999.

[18] L. Quiming, "Scenario Generator for Manets," available from http://www.comp.nus.edu.sg/~liqiming/fyp/scengen/index.html, Apr. 2001.

[19] E. Royer and C. Perkins, "Multicast Operation of the Ad Hoc On-Demand Distance Vector Routing Protocol," *Proc. ACM Mobicom '99 Conf.,* pp. 207-218, Aug. 1999.

[20] P. Sinha, R. Sivakumar, and V. Bharghavan, "MCEDAR: Multicast Core Extraction Distributed Ad Hoc Routing," *Proc. Wireless Comm. and Networking Conf.,* 1999.

[21] 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, 1998.

[22] J. Yoon, M. Liu, and B. Noble, "Random Waypoint Considered Harmful," *Proc. IEEE/INFOCOM Conf.,* 2003.

Kumar Viswanath received the PhD degree in Computer Engineering from the University of California, Santa Cruz and MS degree in electrical engineering from the University of Florida, Gainesville. He has been active in the field of internetworking and group communications in ad hoc networks since 1999. Over the years, his research interests have included distributed sytems, communications, and signal processing and networking protocols for ad hoc networks. He is currently employed as a Research Scientist with NTT Multimedia Communications Labs at Palo Alto.

Katia Obraczka received the BS and MS degrees in electrical and computer engineering from the Federal University of Rio de Janeiro, Brazil, and the MS and PhD degrees in computer science from the University of Southern California (USC). She is an associate professor of computer engineering at the University of California, Santa Cruz (UCSC). Before joining UCSC, she held a research scientist position at USC's Information Sciences Institute and a research faculty appointment at USC's Computer Science Department. Her research interests include computer networks, more specifically, network protocol design and evaluation in wireline as well as wireless (in particular, multihop ad hoc) networks, distributed systems, and Internet information systems. She is a member of the IEEE.

Gene Tsudik received the PhD degree in computer science from the University of Southern California in 1991. He is now a professor in the Computer Science Department at the University of California, Irvine. He has been active in the area of internetworking, network security, and applied cryptography since 1987. Before joining the University of California, Irvine, in 2000, he was a project leader at IBM Research, Zurich Laboratory (1991-1996) and the USC Information Science Institute (1996-2000). Over the years, his research interests included: internetwork routing, firewalls, authentication, mobile network security, secure e-commerce, anonymity, secure group communication, digital signatures, key management, ad hoc network routing, and, more recently, database privacy and secure storage. He has more than 90 refereed publications and patents. He is currently serving as the associate dean of Research and Graduate Studies in the School of Information and Computer Science at the University of California, Irvine.