Interoperability of Multicast Routing Protocols in Wireless Ad-Hoc Networks

Kumar Viswanath and Katia Obraczka
University of California, Santa Cruz
Santa Cruz, CA 95064

January 20, 2004

Abstract

Although MANETs have typically been considered as isolated, stand-alone networks with no connection to the Internet, it is quite likely that future internetworks may interconnect numerous MANET clouds, each running different routing mechanisms. This could be primarily dictated by administrative constraints or different network functional and performance requirements and can result in scenarios where nodes belonging to different MANET clouds running different routing mechanisms may wish to communicate with one another. In this paper we introduce techniques to allow interoperability of various multicast routing protocols in MANETs. In particular we investigate two different interoperability techniques i.e flooding based interoperability and facilitator assisted interoperability. Our preliminary evaluation of the two techniques indicates that the flooding based mechanism is best suited for scenarios involving intermittent communications between different ad-hoc domains whereas the facilitator assisted approach is best suited for data intensive applications such as video-conferencing.

Key Words: Ad-Hoc Networks, Interoperability, Mobile Computing, Multicast, Routing protocols

1 Introduction

Mobile multi-hop ad-hoc networks (MANETs) are self-organizing networks which do not rely on a fixed 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.

Self organizing MANETs have been gaining momentum in recent years on account of capabilities such as ease of deployment, unrestricted mobility, self-organization and management. These capabilities make MANETs very attractive to applications ranging from simple wireless network coverage in offices and warehouses to emergency rescue/response and military operations, law enforcement, and smart environments (including intelligent transport systems). More recently, group-oriented services, such as impromptu conferencing and data dissemination, have been identified as one of the primary classes of applications targeted by MANETs and, as a result, a number of multicast routing protocols have been proposed [7, 19, 6, 23, 27, 8]. These protocols can be classified according to two different criteria, namely (1) how they maintain routing state and (2) what routing structure they build to forward data packets. According to the first criterion protocols are classified as proactive or reactive. Proactive protocols maintain routing state throughout their operation, while reactive mechanisms reduce the impact of frequent topology changes by acquiring and maintaining routes on demand. Using the second criterion, existing protocols are either tree- or mesh-based.

Although current multicast routing mechanisms perform well in constrained mobility MANETs, it has been shown ([10, 14]) that their performance degrades under more stringent network conditions of mobility and traffic load. In general, we believe that no single routing (either unicast or multicast) protocol is optimal for all MANET scenarios given the diverse nature and wide range of operating conditions. As a result, different routing mechanisms will be employed depending on the functional and performance requirements (reliability, overhead, load balancing, security, etc.) and operating conditions (mobility, loss pattern, traffic load, etc.).

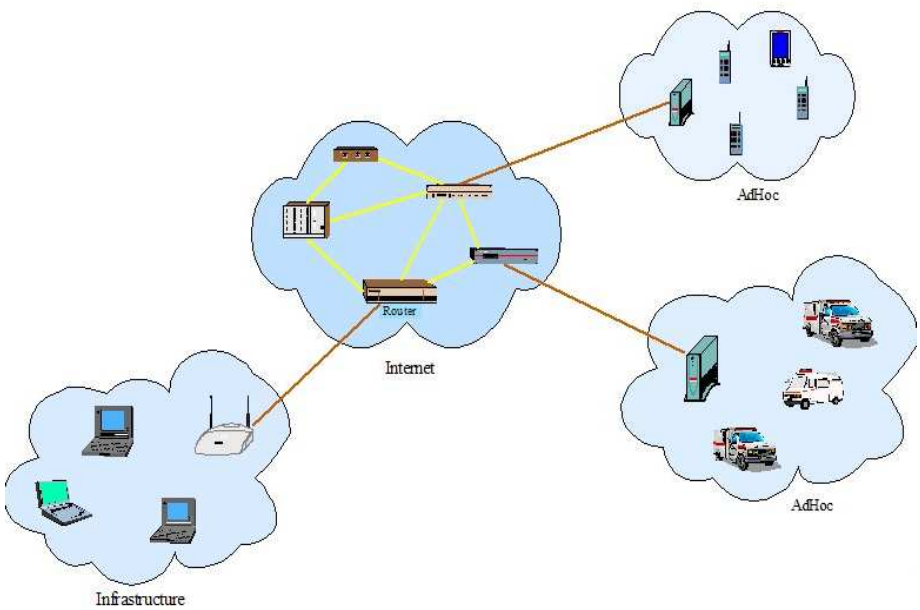

Figure 1: Internetwork with leaf infrastructure- and ad-hoc networks.

Typically, MANETs have always been considered to be isolated, stand-alone networks with no connection to the Internet. However, we envision that future internetworks will consist of a wired backbone and a collection of wired, fixed-infrastructure mobile, and ad hoc networks as leaves as illustrated in Figure 1. We believe that a "global" routing solution for future internets will include specialized solutions for each type of network, as well as mechanisms for integrating these solutions. The choice of routing mechanism could be primarily dictated by administrative constraints, application requirements, operating conditions, or even by varying implementations available from network providers. For example, commercial deployments available from NovaRoam [13] use AODV [17] and TORA [16] as routing protocols, whereas MANET products available from Firetide [5] use TBRPF [15] as the routing mechanism.

Based on these observations, our premise is that, in future internets, different routing mechanisms will coexist and thus calls for mechanisms that allow different protocols to interoperate so that nodes in different *routing domains* can communicate. Our main focus is on multicast communication, in particular, enabling a single multicast group span different MANET clouds and consequently, routing domains.

Further, mobility may cause nodes to migrate from one cloud to another. Node mobility not only raises interoperability and integration issues, but also calls for mechanisms to actively, switch on-the-fly among different multicast routing mechanisms as a mobile node moves across MANET clouds.

In this paper, we focus primarily on interoperability issues in internets consisting of several MANET clouds running different multicast routing mechanisms. To our knowledge, there is little or no experience in the wireless network research community in multicast (or unicast) routing protocol interoperation or adaptation (although, some proposals have been floated in the IETF in the context of wired networks [4, 24, 11]). We introduce two different interoperability mechanisms and evaluate their effectiveness for different MANET environments, including scenarios motivated by typical MANET applications. In particular we investigate flooding based interoperability and facilitator assisted interoperability approaches. The initial simulation results indicate that the flooding based interoperability mechanism is best suited for scenarios involving infrequent communication between different routing domains and when group membership is dense due to flooding's inherent redundancy whereas the facilitator assisted approach is best suited for data intensive applications such as video-conferencing.

The rest of the paper is organized as follows. In Section 2, we highlight the important issues and design considerations for interoperability of routing mechanisms in MANETs. Section 3 overviews our interoperability approaches and Section 4 describes the implementation details. Section 5, describes the simulation environment used, including a description of the simulation parameters. In Section 6, we present simulation results showing how the proposed interoperability techniques perform under various MANET scenarios. Section 7 highlights related work efforts and Section 8 present concluding remarks as well as items for future work.

2 Design Goals

Analogous to wired internets, we define *routing domain* (or *cloud*) as a collection of MANETs under a single administrative control running the same routing protocol. When designing the proposed interoperability mechanisms, the main goals to address include:

- **Scalability**: state that needs to be maintained at nodes and exchanged across domains to allow interoperation should be minimal and scalable with the number of nodes.
- **Prevention of routing loops**: Assuming that the intra-domain routing protocol is loop-free, the interoperabilty mechanism should ensure that the routes across domains are loop free as well.
- **Generality**: The interoperability mechanism should be allow interoperation of various multicast routing techniques. This not only includes interoperability **among** different mesh- (e.g., [7], [2] [6]) and tree-based protocols (e.g., [19], [27]), but also interaction **between** mesh- and tree-based routing. In other words, a domain running a mesh-based protocol should be able to communicate with a domain running a tree-based protocol, and vice-versa.
- **Minimal impact on intra-domain routing**: Each routing domain should be able to choose the routing technique best suited to meet the requirements of its driving applications and network conditions, independently of other domains. Further, the interoperability mechanism should have a minimal effect on the operation of intra-domain routing protocols.

3 Interoperability Mechanisms

In this section we provide an overview of the proposed interoperability mechanisms. Some of the principles that guided our design, which are outlined in Section 2 above, are often conflicting. In particular, if the design favors protocol efficiency and scalability, it will sacrifice protocol generality. This tradeoff led to the development of two interoperability mechanisms, namely (1) flooding-based and (2) facilitator-assisted interoperability. We review these two techniques below.

3.1 Flooding-Based Interoperability

The basic premise of this approach is to use simple flooding to route data across routing domains. The main advantage of such flooding-based interoperability is its simplicity and the fact that it requires no cross-domain route establishment. Furthermore, it is quite easy to prevent routing loops (usually by adding sequence number and a maximum hop count information to packets). Another key benefit of flooding-based interoperability is that it is simple enough to allow interoperation of most routing protocols. However, the main drawback, which is inherent of flooding-based approaches in general, is that it results in excessive redundant transmissions, sacrificing scalability for generality.

Given these features, flooding-based interoperability is suitable for scenarios where receiver groups are quite dense and communication across domains is intermittent. An example where this approach would be beneficial is in the case of emergency response/natural disaster type scenarios. In such circumstances, the rescue operation may consist of several ad-hoc clouds of fire-fighters, medical, search and recover personnel, etc. A major part of the communication in these circumstances is expected to be intra-domain with less frequent inter-domain exchanges (e.g., to obtain feedback and situation appraisal from the co-ordinating stations). Considering that inter-domain communications are most likely to be short-lived, the expense of cross domain route establishment and maintenance (which is the approach taken by the other proposed interoperability scheme) may be unjustified.

3.2 Facilitator-Assisted Interoperability

This approach to interoperability is motivated by scenarios involving frequent inter-domain communications such a video conferencing since flooding-based interoperability mechanisms can prove to be quite expensive for such applications. Facilitator-assisted interoperability requires extra functionality to be assigned to a small set of nodes in each domain referred to as *facilitators*. These special nodes are similar in functionality to Multicast Border Routers (MBRs) [24] in wired networks. In wired networks, MBRs are responsible for connecting two multicast routing domains by sharing their forwarding caches. The multicast protocol running on each MBR sends its forwarding table entries to a shared cache. The shared cache functions as the bridge between multicast trees in neighboring domains. Analogous to MBRs, facilitators in each domain act as entry and exit points for all cross-domain communication. They can also function as normal nodes (i.e, sources, multicast receivers or forwarding group members) in the context of their own routing domain. Facilitators in each domain are essentially responsible for creating and maintaining links between neighboring domains through the use of periodic signalling.

Facilitator-assisted interoperability favors scalability since it uses the existing structure created by the underlying routing protocols to forward data, avoiding global flooding. However, in order to do so, it sacrifices generality as it may require changes to the underlying routing protocol.

For example, in sender initiated protocols (e.g., ODMRP), facilitator nodes also have to act as "dummy" senders to establish group membership.

4 Protocol Description

One key modification required by our interoperability mechanisms is the addition of a *composite header* for all routing layer messages. The *composite header* ensures that nodes in different routing domains running different routing mechanisms have some minimal knowledge about cross-domain routing messages. The information carried by the composite header include:

- `Protocol Type` {*proto*}: `Specifies the type of the protocol e.g MAODV, ODMRP, AMRIS etc.`
- `Message Type` {*mtype*}: `Specifies the type of message e.g request, reply, data`
- `Source Address` {*srcAddr*}: `Address of node initiating the message`
- `Destination Address` {*dstAddr*}: `Address of destination node or multicast group`
- `Next Hop` {*nextAddr*}: `Next hop towards destination`
- `Sequence Number` {*seq_number*}: `Required to detect duplicate messages`

The use of the different fields will become clear in the following sections as we illustrate the operation of the interoperability mechanisms through examples.

4.1 Flooding-Based Interoperability

In the flooding based approach, when nodes receive routing messages, they first check the composite header for the *proto* field. If protocol types are the same, the packet is handled by the default routing operation. However, if the protocol types are different, then nodes check the *mtype* field to determine the message type. All routing messages (i.e., route requests, replies etc.) are dropped silently. However, if the *mtype* field indicates a data packet, then the node re-broadcasts the packet to its immediate neighbors. Nodes also cache packet headers to prevent looping and unnecessary retransmissions. The basic operation of the protocol is illustrated through example below.

Figure 2 illustrates the initial route setup across two different MANET clouds using flooding-based interoperability. The domain on the right runs a tree-based multicast routing protocol while the left domain employs a mesh-based protocol. When nodes receive routing messages from neighboring domains these routing messages are not forwarded and dropped silently.

Data forwarding across domains is illustrated in Figure 3. In this particular example, the source is in the tree domain. Data packets are handled normally in the tree domain in accordance with the underlying routing protocol. However, when the packets are received in the mesh domain, nodes first add the relevant information to their packet caches and broadcast the data to their immediate neighbors. Each node that receives the data packet essentially does a one-hop broadcast after recording the packet header in its cache leading to flooding of data in the mesh domain.

The pseudo-code corresponding to our implementation of flooding-based interoperability is shown below.

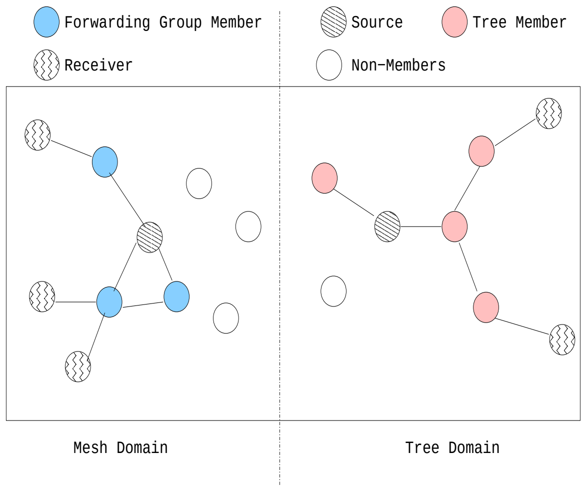

Figure 2: Flooding-based interoperability: initial route setup across two routing domains.

```
case: Non Member of Multicast Group
 if  (proto != my_proto) {
    if (mtype == data) {
       Add packet header to cache
       Rebroadcast Data Packet
      }
    else if (mtype != data) {
      drop packet
     }
  }

case: Member of Multicast Group
 if  (proto != my_proto) {
   if  (mtype == data) {
       Add packet header to cache
       Send packet to transport layer
       Rebroadcast Data Packet
      }
    else if (mtype != data) {
           drop packet
      }
 }
```

Using a variation of the simple flooding approach, it may be possible to further reduce the number of redundant retransmissions. The idea is to use *scoped flooding* [26] which employs different heuristics to decide whether a packet should be re-broadcast. 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. When a node receives a data packet 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, it does not re-broadcast the packet. The assumption is that if there is significant overlap in the neighborhood of the two

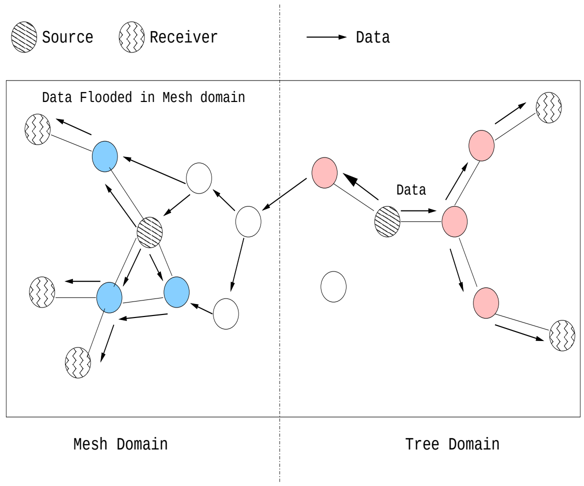

Figure 3: Flooding of data in mesh routing domain

nodes, it is likely that the packet has already been received by the neighbors.

4.2 Facilitator-Assisted Interoperability

As mentioned in Section 3.2, facilitators in each domain maintain links to facilitators in neighboring domains through a signaling protocol consisting of two types of messages: (1) *fac-request* and *fac-reply*. *Fac-requests* are periodically initiated by all facilitators for inter-domain link creation and maintenance. Facilitators of neighboring domains respond to *Fac-requests* by sending *Fac-replies* towards the source of the request.

The following example illustrates the sequence of operations and protocol interactions involved in facilitator-assisted interoperability.

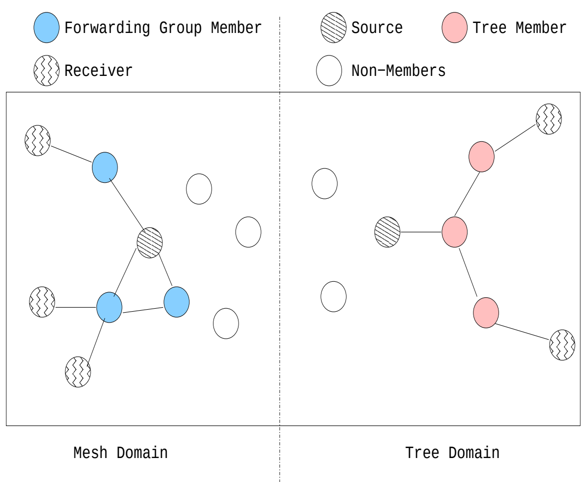

Figure 4: Facilitator-assisted interoperability: initial route setup across two routing domains

Figure 4 depicts a scenario in which a multicast group spans two different routing domains, each running a different multicast routing protocol. More specifically, the domain on the right employs a tree-based routing protocol while the domain on the left a mesh-based protocol. The initial route setup in each domain is shown in Figure 4.

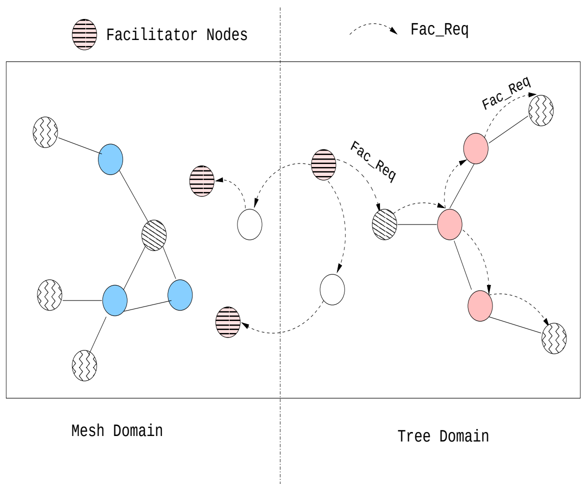

Figure 5: Facilitator request

As shown in Figure 5, certain nodes in each domain are selected as facilitators [1]. For the purpose of this discussion, we assume that facilitators are randomly selected. When a node is selected to be a facilitator for a domain, it first checks whether it is a member of the multicast group. If it is not a member, it sends out a join group message which is specific to the underlying multicast protocol used in its domain. This is required to ensure that all facilitators have routes to multicast groups in their own domain. As shown in Figure 5, a facilitator in the tree domain broadcasts a *fac-request* message. The composite header for this message is of the form `< proto_type tree, msg_type FAC_REQ, srcAddr, mcastAddr, seqNumber #1 >`. When nodes in the tree domain receive the *fac-request*, they add the request to their *fac request table*, set up reverse routes to the source of the request, and forward the request. Similarly, when nodes in the mesh domain receive the *fac-request*, they add the request to their *fac request table*, set up reverse routes to the source of the request, and forward the request. Thus the *fac-request* is propagated till it is received by a facilitator in the mesh domain. As shown in Figure 6, after ensuring that the request is not a duplicate, the facilitator in the mesh-domain sends out a *fac-reply* to the next-hop towards the source of the request. The composite header for this message is: `< proto_type mesh, msg_type FAC_REP, srcAddr, mcastAddr, next-hop, seqNumber #1 >`. The next-hop, on receiving the *fac-reply*, after ensuring that the reply is not a duplicate, checks to see if it is already a member of the mesh. If it is not a mesh member, it becomes one and forwards the packet along the reverse route towards the source of the *fac-request*. All intermediate nodes along the path to the source in both tree and mesh domains are incorporated as members for the tree or mesh. The forwarding group status for the incorporated nodes is periodically refreshed by the *fac-requests*. The effect of the *fac-request* - *fac-reply* exchange is to create links between facilitators in each domain.

Fig 7 depicts the traversal of data across neighboring domains. In each domain, data is forwarded according to the underlying protocol along the existing tree or mesh. The only difference is that certain nodes that are incorporated as part of the tree or mesh due to the *fac-request* -

[1] We discuss facilitator selection in greater detail in Section 4.3.

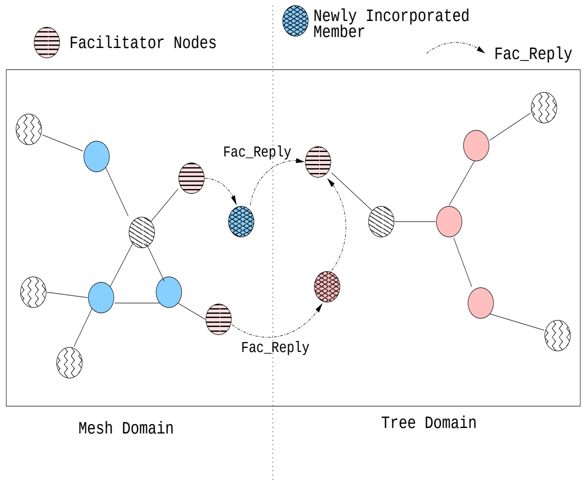

Figure 6: Facilitator reply

fac-reply exchange are also responsible for forwarding data.

The pseudo-code corresponding to our implementation of facilitator-assisted interoperability is shown in figure 8.

4.3 Facilitator Selection

As explained above facilitators are responsible for enabling cross-domain communications by creating and maintaining mesh links across domains. Hence, the choice of facilitators can impact the performance of group communications both within and across domains. When selecting facilitators, one of the essential requirements is that all facilitators be members of the multicast group. In receiver-initiated multicast routing (e.g., MAODV), this ensures that facilitators have routes to the multicast group within their own domain. However, in case of sender-initiated routing protocols (e.g., ODMRP), facilitators also act as "dummy" senders for the group to ensure that they have paths to multicast receivers within their own domain.

In scenarios where MANET clouds are wireless extensions of the internet, certain nodes are typically assigned special functionality (e.g., gateways, DNS service providers, proxies, web caches, etc.). These nodes naturally lend themselves to be elected as facilitators. The number of facilitators in each domain is another important factor in determining performance. In scenarios where nodes are mobile, it is necessary that each domain has sufficient number of facilitators to ensure redundant links to neighboring domains. However, increasing number of facilitators can also potentially degrade performance on account of the increase in facilitator-related control messages. The number of facilitators should be typically assigned based on the requirements of the network in terms of control overhead and data delivery guarantees. We have conducted a preliminary study of the impact of facilitator selection on the performance of facilitator-assisted interoperability and present our results in Section 6.

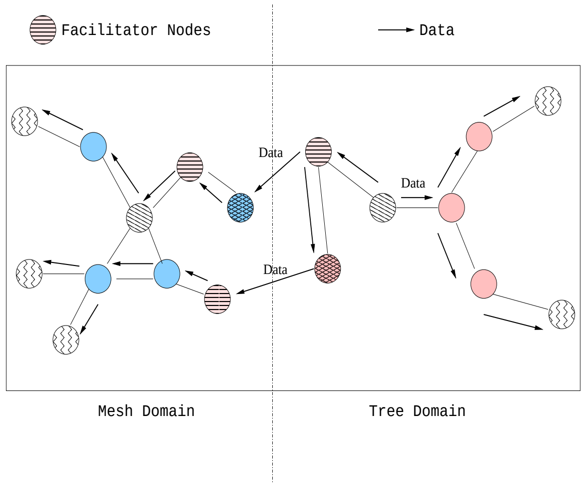

Figure 7: Forwarding data

5 Simulation Setup

We used `Qualnet` [20] as the simulation platform. The simulation setup essentially consisted of two routing domains, one running a tree-based routing protocol and the other a mesh-based protocol. MAODV [19] was chosen as representative of tree-based routing while ODMRP [7] was chosen to represent mesh-based routing techniques. It should be noted that MAODV is a receiver initiated protocol, i.e, receivers send explicit join messages to graft themselves to the multicast tree. However, ODMRP is sender initiated, i.e, the multicast mesh is formed by join queries transmitted by senders prior to data transmission. Thus, in the ODMRP domain, facilitators had to function as "dummy" senders to ensure that they had routes to multicast groups in their own domain.

We use two type of MANET scenarios in our simulations. In "synthetic" scenarios, parameters such as mobility, number of facilitators, traffic sources, and number of multicast receivers are varied over an arbitrary range of values. We also define more "concrete" environments reflecting specific MANET applications, namely exhibition/symposium scenario. The exhibition/symposium scenario was generated using the scenario generator presented in [18] and is described in detail in section 6.3.

In our simulations, each domain consists of 50 nodes randomly placed in a 1000 m^2 field. The domains had sufficient overlap to enable communications from one domain to another. Nodes' channel bandwidth was set to 11 Mbit/sec and their transmission range is 225 meters. All member nodes join at the start of the simulation and remain members throughout the experiment.

Random waypoint was used to model node mobility nodes. For our simulations, nodes were restricted to move only within their own domains. Each source had a constant bit rate (CBR) traffic generator generating 5Kbps. The data payload size was fixed at 256 bytes.

We use the following metrics in evaluating the performance of the different interoperability mechanisms.

- **Packet delivery ratio** is computed as the ratio of total number of packets received by the nodes to the total number of packets transmitted 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`

```
case: Facilitator Node

Proc Init-Fac-Node:
if (!member of multicast group) {
   Join_Group
  }
  Send Periodic Fac_Request

Proc Fac-Node-Recv-Packet:
if (proto != my_proto) {
  if (mtype == FAC-REQ) {
      Add Req to Fac Req Table
      Send-Fac-Reply
     }
  if (mtype == FAC-REP) {
     if (!initiator of Request)
       forward along reverse path to source
        Add Rep to Fac Rep Table
     }
  if (mtype == DATA) {
      One hop broadcast to neighbors
     }
}

if (proto == my_proto) {
  if (mtype == FAC-REQ) {
     Add Req to Fac Req Table
     Set Reverse Path to source
     Forward Fac Req
    }
  if (mtype == FAC-REP) {
     Add Rep to Fac Rep Table
     Forward along reverse path to source
    }
  if (mtype == DATA) {
     One hop broadcast to neighbors
    }
}
```

```
case: Regular Nodes

Proc Recv-Packet:
if (proto != my_proto) {
  if (mtype == FAC-REQ) {
     Add Req to Fac Req Table
     Set Reverse Path to Source of Req
     Forward Fac Req
   }
  if (mtype == FAC-REP) {
     if (nextAddr == my_addr) {
       if (!member of group)
          Set as Forwarding Group Member
       Add to Fac Rep Table
       Forward along reverse path to Source
      }
   }
  if (mtype == DATA) {
     if (forwarding member  of group)
          One hop broadcast to neighbors
       else drop packet
   }
}
if (proto == my_proto) {
  if (mtype == FAC-REQ) {
      Add Req to Fac Req Table
      Set Reverse Path to source of Req
      Forward Fac Req
   }
  if (mtype == FAC-REP) {
     if (nextAddr == my_addr) {
        if (!member of group)
            Set as Forwarding Group Member
         Add to Fac Rep Table
    Forward along reverse path
      }
  }
 if (mtype == DATA) {
     if (forwarding member  of group)
         One hop broadcast to neighbors
      else drop packet
     }
}
```

Figure 8: Pseudo Code for Facilitator Assisted Interoperability

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.

For the flooding-based interoperability approach, in addition to the underlying routing protocol overhead, control overhead also includes all data header bytes forwarded by network nodes and overhead due to the `composite header`.

In the case of facilitator-based interoperability, the routing overhead includes `fac requests`, `fac replies` and also accounts for the `composite header` in the data forwarding process.

6 Simulation Results

In this section, we report simulation results comparing the different interoperability approaches. We ran each simulation (keeping all parameters constant) five times, each time using different seed values. Each data point in the graphs below, represents the average across five runs.

6.1 Effect of Mobility

For the mobility experiment, 10 nodes from the tree-domain were randomly chosen as traffic sources. Each source transmitted 5Kbps and thus the overall network load was 50Kbps.The multicast group comprised of 20 receivers in each domain. Further, 5 nodes from each domain were randomly assigned as facilitators for the facilitator-based approach. Average node speed was varied from 3.6 to 100 kms/hr.

Figure 9 depicts the reliability performance of the two interoperability approaches as a function of node mobility.

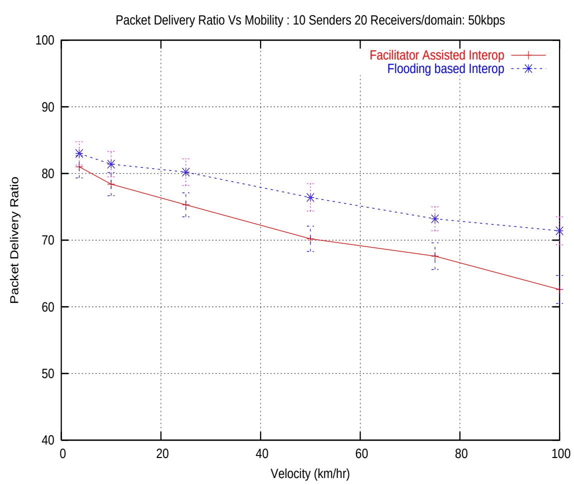

Figure 9: Packet delivery ratio as a function of node mobility

We observe that flooding-based interoperability exhibits higher reliability than the facilitator assisted approach. The difference in delivery ratio is around 10% at 100km/hr. This trend is quite intuitive since flooding based approaches are more resilient to link failures. In the facilitator assisted approach, node mobility may cause links between facilitators in the mesh and tree domain to frequently go down. Hence no cross-domain communication can occur until the links are refreshed by new `fac requests`. Further it can be seen that delivery ratio for both approaches decreases with increase in mobility. This decline in the delivery ratio can be attributed to the behavior of the underlying routing protocol. Increased mobility causes frequent link changes and requires MAODV to reconfigure the multicast tree more frequently to prevent stale routing information. A number of data packets in the tree domain are dropped due to route failures and hence never transmitted to the mesh domain.

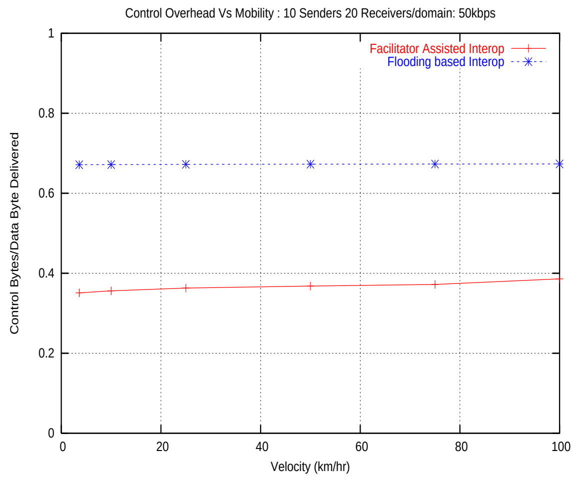

Figure 10: Control overhead as a function of node mobility

Figure 10 plots control overhead per data byte transferred as a function of mobility and shows that control overhead for both approaches increases with node mobility. As average node speed increases, link failures in the tree-domain triggers route-error messages and causes frequent tree configuration. This increases the flow of control messages thus increasing the overhead. However, the key result from the figure is that the routing overhead of the flooding based approach is almost twice that of the facilitator approach. As mentioned earlier, the flooding based approach results in a large number of redundant transmissions in the mesh-domain which considerably increases the overhead.

6.2 Facilitator Selection

As mentioned in section 4.3 the selection of facilitators can significantly affect the performance of group communications both within and across domains. Intuitively, we can expect the performance to increase as we increase number of facilitators since the mesh links between facilitators in different domains becomes richer. This provides greater redundancy and resilience to link failures and node mobility. In this set of experiments 10 randomly selected nodes act as sources generating CBR traffic at a rate of 5Kbps. The number of multicast receivers in each domain was increased to 40 and node mobility was fixed at 10 kms/hr. The number of facilitators is varied from 5-40 and we observe its impact on the performance metrics.

Figure 11 shows the impact of number of facilitators on packet delivery ratio.

As discussed earlier packet delivery ratio increases as number of facilitators increase. However when the number of facilitator is greater than 20 the delivery ratio starts decreasing. For this

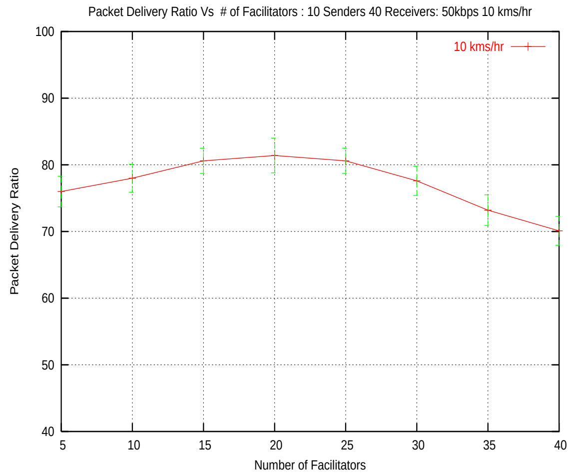

Figure 11: Packet delivery ratio versus number of facilitators

particular setup, increasing the number of facilitators up to 20 strengthens the connectivity between facilitators in the tree and mesh domain. This makes the interoperability mechanism more resilient to link failures. However, increasing the number of facilitators also has the negative effect of increasing the control overhead due to greater cross-domain facilitator exchanges. Data packets have to contend with facilitator control messages which results in greater channel contention and packet drops due to collisions.

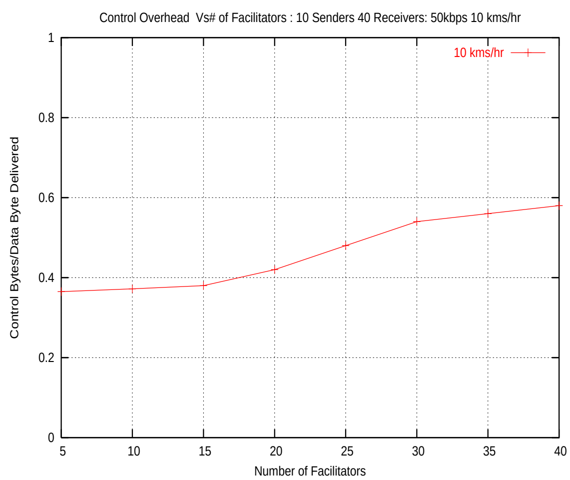

Figure 12: Control overhead versus number of facilitators

Figure 12 shows the impact of number of facilitators on routing overhead. As expected the control overhead increases with increase in number of facilitators. One subtle side-effect is that a larger number of non-member nodes in the tree and mesh domains are also incorporated as members to forward data due to facilitator exchanges. Although this increases redundancy it also results in larger number of redundant data retransmissions.

We have also investigated the impact of mobility and traffic load on the facilitator selection process. In these experiments, 10 nodes were chosen as traffic sources and 20 nodes in each domain were chosen as multicast receivers. Node mobility and traffic load was varied over arbitrary values and the impact on the overall performance is depicted in the graphs below.

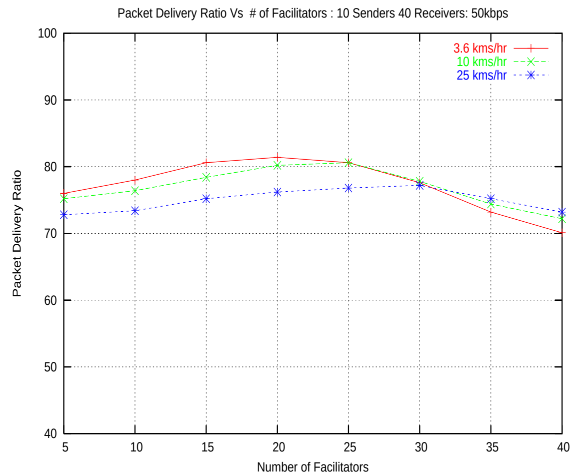

Figure 13: Packet delivery ratio versus number of facilitators for different node mobility

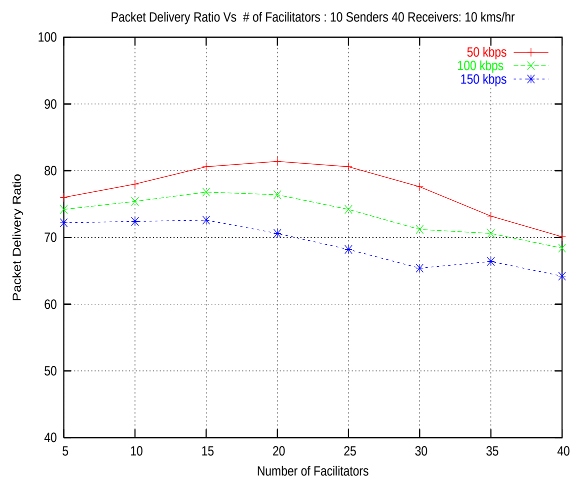

Figure 14: Packet delivery ratio versus number of facilitators with varying traffic load

Figures 13 and 14 indicate that the optimum value of the facilitators for a particular scenario is dependent on node mobility and traffic load. From Figure 13 it is observed that as node mobility increases, the interdomain links between facilitators are more prone to failures. In high mobility scenarios, increasing the number of facilitators increases the overall packet delivery ratio by providing higher number of redundant links. Similarly, Figure 14 indicates that for a traffic load of 50 kbps the optimum number of fcilitators is around 20. However, the optimum number of facilitators decreases to 15 as traffic load is increased to 150 kbps. As noted earlier, increasing number of facilitators increases the control traffic. It also has the side-effect of incorporating a larger number of non-member nodes to participate in the data forwarding process which results in higher packet contention and data loss due to collisions.

One key observation from the above experiments, is that the performance of the interoperability mechanism is heavily dependent on the facilitator selection criteria. This requires that facilitors be chosen adaptively, based on the underlying network conditions. One adaptive criteria for chosing facilitators can be based on node mobility. In this method, all nodes monitor the relative mobility of their one-hop neighborhood through periodic hello messages. If the relative mobility of the one-hop neighborhood is below a certain threshold (over a time window) then the node can choose to elect itself as a facilitator. This ensures relatively robust intermediate links between the facilitators across domains increasing the overall performance of the interoperability mechanism. We are currently investigating other adaptive algorithms for the facilitator selection process.

6.3 Exhibition/Symposium Scenario

In this experiment we evaluate the performance of flooding based and facilitator assisted interoperability in a realistic MANET scenario. The exhibition/symposium scenario was generated using the `scengen` [18] tool. The scenario consisted of 75 nodes in a 1000 m^2 field in each domain. The nodes were divided into the following categories: 20 stationary *exhibition booths*, 20 *audience* nodes and 35 *wanderer* nodes. The exhibition booths were randomly distributed within their respective domains and were assigned as facilitators for the facilitator assisted mechanism. The members of the audience group moved with speeds ranging from 2-5 m/s and their movement was modeled using brownian motion. The movement of the audience group was restricted to a limited area inside their own domains. The wanderer nodes moved according to the random-waypoint model with speeds ranging from 1-5 m/s. Unlike the audience group, wanderers were capable of moving throughout their entire domain. The receiver group consisted of 50 nodes randomly selected in each domain. All exhibition booths in both domains functioned as traffic sources, generating CBR traffic at the rate of 5Kbps.

Table 1 summarizes simulation results for the exhibition scenario.

Exhibition scenario		
Interoperability Mechanism	Delivery ratio %	Routing overhead (Bytes Xmitted/data bytes recvd)
Flooding	80.2	0.57
Facilitator	76.4	0.34

Table 1: Exhibition Scenario

Similar to the results from our prior experiments the flooding based approach had a slightly higher packet delivery ratio as compared to facilitator assisted interoperability mechanism. However, the control overhead was about 60% higher compared to the facilitator based approach.

6.4 Discussion

Our simulation results confirm that flooding based interoperability is more reliable but incurs significantly higher routing overhead. However, in cases where generality is important (e.g pre-deployed MANETs) then flooding based interoperability approaches may be the only choice. In general, flooding based approaches are best suited for infrequent communications between different domains. On the other hand for data intensive multicast applications such a video-conferencing the facilitator assisted approach is better suited. Facilitator assisted approaches require minor modifications to the intra-domain routing protocol and hence trade-off reliability for scalability.

7 Related Work

To our knowledge, there has been very little experience in the wireless network research community in multicast (or unicast) routing protocol interoperation or adaptation. However, interoperability issues in the context of wired networks have been addressed in some detail and several proposals have been floated in the IETF. Hierarchical DVMRP (HDVMRP) [25] which was proposed as an inter-domain routing protocol aims to interconnect multiple domains by flooding data packets to boundary routers. HDVMRP divides the flat routing region into non-overlapping regions. DVMRP [3] is used as the multicast routing protocol within domains and also as the intra-domain routing protocol between different regions. The Border Gateway Multicast Protocol (BGMP) architecture [9] addresses the scaling problems of approaches such as HDVMRP. BGMP consists of two complementary protocols for inter-domain multicasting. The Multicast Address Set Claim (MASC) proposes an hierarchical address allocation scheme for dynamic address allocation to domains. The Border Gateway Multicast Protocol (BGMP) builds bidirectional shared trees across domains without interfering with the intra-domain routing protocols. Hierarchical Multicast Routing (HIP) [22] was also motivated by the scalability problems of DVMRP and provides means for routing across heterogeneous domains. HIP uses the concept of virtual routers (VRs) to organize all border routers of a domain so that they appear as a single router to the higher level tree. Ordered Core Based Trees [21] is used for inter-domain routing. RFC 2715 [24] addresses interoperability requirements for multicast protocols in wired networks. The proposed framework attempts to provide efficient interoperability among different multicast routing protocols such as DVMRP [3], MOSPF [12], CBT [1] etc.

8 Conclusions and Future Work

In this paper, we introduce interoperability techniques to facilitate seamless multicast communication between nodes spanning heterogeneous domains. In particular we propose two different interoperability techniques, i.e, flooding-based and facilitator assisted interoperability. The flooding-based interoperability technique has the advantage of being simple in terms of implementation and requires no explicit cross-domain route establishment protocol. This is beneficial in situations where it may be difficult to change the existing network infrastructure but interoperability is still desired. The facilitator based approach, on the other hand, requires the addition of special functionality to a small subset of nodes in each domain. The functionality of facilitators is quite similar to that of Multicast Border Router's (MBRs) in wired domains. This approach is well suited for scenarios involving frequent inter-domain communications such a video conferencing since flooding-based interoperability mechanisms can prove to be quite expensive for such applications. However, the side-effect of flooding-based approaches is better reliability on account of the redundant transmissions. The facilitator based approach on the other hand favors scalability as opposed to generality since it requires minor modifications to the underlying routing protocol behavior.

Given that, interoperability of routing protocols is inevitable, both flooding as well as facilitator assisted approaches need to be researched further. We are currently investigating other flooding based techniques such a *scoped flooding* in which nodes use certain heuristics to determine whether to rebroadcast data to neighboring domains. Another area for further research is to study the impact of facilitator selection and to research algorithms for optimizing the facilitator selection criteria. We also plan to study the feasibility of extending our proposed approaches for unicast protocol interoperation.

References

[1] A.J. Ballardi, P.F. Francis, and J. Crowcroft. Core based trees (cbt). *Proc. of the ACM SIGCOMM'93*, August 1993.

[2] Ching-Chuan Chiang, Mario Gerla, and Lixia Zhang. Forwarding group multicast protocol (FGMP) for multihop, mobile wireless networks. *Cluster Computing*, 1(2):187–196, 1998.

[3] S.E. Deering and D.R. Cheriton. Multicast routing in datagram internetworks and extended lans. *ACM Transactions on Computer Systems*, 8(2):85–110, May 1990.

[4] D. Estrin et. al. Pim border router specification for connecting pim-sm domains to a dvmrp backbone. Internet draft draft-ietf-idmr-PIM-SM-spec-09.ps, October 1996.

[5] FireTide. http://www.firetide.com.

[6] J. Garcia-Luna-Aceves and E. Madruga. A multicast routing protocol for ad-hoc networks. In *Proc. of INFOCOM'99*, pages 784–792, March 1999.

[7] 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/.

[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] Satish Kumar, Pavlin Radoslavov, David Thaler, Cengiz Alaettinoglu, Deborah Estrin, and Mark Handley. The masc/bgmp architecture for inter-domain multicast routing. In *SIGCOMM*, pages 93–104, Vancouver, British Columbia, Sep 1998. ACM.

[10] S.J. Lee, W. Su, J. Hsu, M. Gerla, and R. Bagrodia. A performance comparison study of ad hoc wireless multicast protocols. *In Proceedings of the IEEE Infocom 2000*, March 2000.

[11] D. Meyer. Some issues for an inter-domain multicast routing protocols. Internet draft draft-ietf-mboned-issues-03.txt, November 1997.

[12] J. Moy. Multicast extension to OSPF. Internet Draft, September 1992.

[13] NovaRoam. http://www.novaroam.com.

[14] Katia Obrazcka, Gene Tsudik, and Kumar Viswanath. Pushing the limits of multicast in ad hoc networks. technical report, 2000.

[15] R. Ogier, F. Templin, and B. Bellur. Topology broadcast based on reverse-path forwarding (tbrpf). IETF manet (draft-ogier-manet-tbrpf-00.txt), 1998.

[16] V.D. Park and M.S. Corson. A highly adaptive distributed routing algorithm for mobile wireless networks. *Proceedings of the IEEE Infocom 1997*, April 1997.

[17] C.E. Perkins and E.M. Royer. Ad-hoc on demand distance vector routing. *Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, pages 90–100, February 1999.

[18] Li Quiming. Scenario generator for manets. Available from http://www.comp.nus.edu.sg/ liqiming/fyp/scengen/index.html, April 2001.

[19] E. Royer and C. Perkins. Multicast operation of the ad-hoc on-demand distance vector routing protocol. *Proceedings of the ACM Mobicom'99*, pages 207–218, August 1999.

[20] Scalable Network Technologies. *Qualnet Wireless Simulator*. Available from http://www.qualnet.com.

[21] Clay Shields and J. J. Garcia-Luna-Aceves. The ordered core based tree protocol. In *INFOCOM*, Kobe, Japan, April 1997. IEEE.

[22] Clay Shields and J. J. Garcia-Luna-Aceves. The HIP protocol for hierarchical multicast routing. In *SIGACT*, pages 257–266, Puerto Vallarta, Mexico, June 1998. ACM.

[23] 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.

[24] D. Thaler. Interoperability rules for multicast routing protocols. Internet draft draft-thaler-interop-00.ps, November 1996.

[25] A. Thyagarajan and S. Deering. Hierarchical distance vector multicast routing for the mbone. *In Proc. of the ACM SIGCOMM*, August 1995.

[26] Kumar Viswanath, Katia Obrazcka, and Gene Tsudik. An adaptive approach to group communications in multi-hop ad hoc networks. *IEEE, International Conference on Networking*, pages 622–633, 2002.

[27] 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.