Performance Comparison of Two Interdomain Routing Schemes for Disruption Tolerant Networks

Mooi Choo Chuah, *Senior Member, IEEE*, Peng Yang, *Member, IEEE*

Lehigh CSE Department
chuah@cse.lehigh.edu, pey204@lehigh.edu

Abstract— **Much work has been done on designing routing protocols for mobile ad hoc networks. However, existing solutions assume that an end-to-end path exists from a source to a destination. Some ad hoc network scenarios e.g. bus-based vehicular adhoc networks are characterized by frequent partitions and intermittent connectivity. Hence, new routing schemes designed for such networks have emerged in the past two years. These disruption tolerant network (DTN) routing schemes assume a flat architecture. However, in real life scenarios, nodes may be from different administrative domains and hence form different clusters. Thus, other means need to be provided to deliver interdomain traffic. In this paper, we first describe two interdomain routing protocols for disruption tolerant networks, namely the gateway-based and the ferry-based approaches. Then, we demonstrated via simulation studies that both interdomain routing schemes provide better delivery performance than the flat routing approach when different groups are segregated from one another. We also show that the ferry-based scheme achieves higher delivery ratio and lower average end-to-end delay for the interdomain traffic than GBIR when the ferry speed is higher than the average node speed. In addition, we demonstrated that the choice of intradomain routing scheme affects the delivery performance of intergroup messages. Using a multihop intradomain routing scheme is better than using a two-hop relay routing scheme. Subsequently, we evaluate how mobility models, number of groups affect the delivery performance of the ferry-based scheme.**

1 Introduction

Wireless ad hoc networks are networks that can be formed dynamically by mobile hosts without any preinstalled infrastructure. Much work has been done in the past to design flat routing schemes for ad hoc networks [1],[2]. The flat ad hoc routing scheme is proven to have poor scalability [3]. Hence, hierarchical routing solutions have been recently proposed. For example, in [4], the authors propose a two level hierarchical ad hoc network where some "backbone" nodes are assumed to have an additional powerful radio to establish long range wireless links among themselves, thus forming a mobile backbone. The backbone nodes are also moving and hence they form yet another ad hoc network. The local subnets can run one routing protocol while the mobile backbone runs another routing protocol.

There are very few papers that address interdomain routing problems for the communication between various ad hoc groups from different administrative domains and possibly with different network configurations. In [5], we designed and compared three gateway-based interdomain routing schemes for mobile ad hoc networks. In [5], we still assume that an end-to-end path between a source and a destination from different domains exists. However, there are some real-life scenarios e.g. battlefields, disaster relief efforts or bus-based vehicular adhoc networks [24] where we may have clusters of nodes that are far away from one another and hence an end-to-end path may not exist either within each cluster of nodes or between clusters.

Recently, a new network architecture [6] called the Disruption Tolerant Network (DTN) has been proposed to allow partitioned nodes or clusters of nodes to communicate with one another. Recent research interests in this area include network architecture design [6],[7], and different routing algorithms for DTNs [8][9][10],[11],[12]. All except [8] are routing schemes designed for delivering intradomain traffic. As mentioned earlier, in real life scenarios, clusters of nodes may form different ad hoc networks running their own intragroup routing schemes. Hence, new interdomain routing schemes need to be designed to allow different clusters to communicate with one another.

In this paper, we assume that special nodes called message ferries may be deployed to provide intergroup delivery services. The routes taken by such ferries can be controlled. We further assume that such ferries can move several times faster than regular nodes. This assumption is quite realistic, e.g., in military scenarios, hum bees moving at 60km/hour can be used as the message ferries while the soldiers (with wireless devices) moving at a fast walking speed of 10km/hour will be regular nodes. With these network assumptions, we focus on understanding how interdomain routing schemes affect the delivery performance of intergroup messages in DTNs. Specifically; we are interested in identifying scenarios where interdomain routing schemes provide better delivery performance than flat routing schemes. In addition, we are also interested in comparing the delivery performance of two interdomain routing schemes for DTNs, namely (a) the gateway-based interdomain routing (GBIR) scheme, and (b) the ferry-based interdomain routing (FBIR) scheme. In GBIR, gateway nodes that can hear other groups are used to deliver intergroup messages while in FBIR scheme, dedicated nodes called message ferries are used to deliver such messages. Intuitively, GBIR scheme will be more useful for scenarios where nodes from different administrative domains may occupy overlapping geographical area while FBIR may be more useful for scenarios where clusters of nodes do not meet and are far away from one another. We are interested in understanding whether there are any scenarios where GBIR will perform better than FBIR.

Our contributions in this paper can be summarized as follows: Via simulation studies, we demonstrate that (a) interdomain routing schemes (aka hierarchical routing approach) provide better intergroup message delivery performance than flat routing schemes when different groups of nodes do not mix together, (b) FBIR performs better than GBIR if different groups are segregated or if the message ferry travels faster than the average node speed when different groups do meet one another, (c) the delivery performance of intergroup messages is affected by the choice of intradomain routing scheme. In addition, we also study how the mobility models, number of groups affect the intergroup message delivery performance of the FBIR scheme. Our results indicate that the UMassBusNet [24] mobility model yields the poorest performance among the four mobility models we considered. Our results with the number of groups indicate that the delivery performance degrades slowly with increasing number of groups. The results also show that the FBIR scheme is scalable.

The remainder of this paper is organized as follows. We provide a brief review of related work in Section 2. In Section 3, we present two interdomain routing schemes in detail. In Section 4, we describe our simulation setup and present our simulation results. We conclude in Section 5 with some discussions on future work.

2 RELATED WORK

2.1 Routing in Intermittently Connected Networks

Several routing schemes have been proposed for DTNs [10],[11],[12],[13],[14]. These different schemes can be grouped into three categories. The first category [9] uses special nodes called ferries to deliver messages between partitioned networks. Ferry routes have significant effect on the data delivery performance, hence they need to be designed efficiently. The second category [11],[12] uses a multihop routing approach where contact history information is used to determine the next hop node to pass a message. For example, in [12], a probabilistic metric called delivery predictability is used to determine if a node needs to pass any stored messages to a new contact that it comes across. The third category [13],[14] uses a two-hop routing approach where the intermediate nodes that receive messages from any source have to store the messages until they can deliver the messages when they come into contact with the destinations of the messages. Sometimes, erasure-coding is used to encode and divide the message into multiple blocks and these different blocks are sent to different relays to increase the chances of a destination receiving a particular message since the destination only needs to receive a certain fraction of the encoded blocks to reconstruct the original message.

2.2 Hierarchical and Interdomain Ad hoc Network Routing

Most of the existing routing researches for ad hoc networks only deal with scenarios where the nodes belong to the same administrative group. A real interdomain routing problem studied so far is the work done in [15],[16] where the authors study how packets can be delivered between ad hoc networks and the wired Internet. Their approaches assume that the nodes wait for a certain period of time for route replies. If no route reply is heard, then the sending node assumes that the destination node is in the wired Internet and proceeds to use Mobile-IP like protocol to register with a foreign agent that can access the wired Internet so that its packets can be delivered to the desired destination outside the ad hoc network.

Our paper [5] is the first that addresses the interdomain routing schemes for communications between different ad hoc groups. In [5], we design and compare three interodmain routing schemes. These schemes are gateway-based schemes where the nodes from one group that can hear nodes from another group will compete to become the gateway for forwarding interdomain route requests. Since the nodes move and hence the gateway nodes may change frequently, the schemes in [5] will not perform well especially when the nodes are sparsely distributed over a large area. The gateway-based schemes described in [5] use limited flooding for gateway selection and gateway information is flooded over the whole network. Such operations result in high overheads. Thus, in this paper, we design a new gateway-based interdomain routing scheme that has better performance than the schemes described in [5]. In [8], the authors propose using message ferries to deliver inter-region messages. Our work differs from theirs in that (a) we also use the message waiting time as a triggering condition for ferry to move to foreign group, and (b) we allow ferry to visit multiple groups before returning. In addition, we evaluate the impact of intra-domain routing protocols on the end-to-end delivery delay of intergroup traffic but the authors in [8] did not. We also investigate when hierarchical interdomain routing schemes will be useful.

3 System Architecture

In this section, we first set boundaries for our work by describing some of the assumptions we make for the system model. Then, we describe the two interdomain routing schemes that we explore in this paper.

3.1 System Model

In this paper, we consider disruption tolerant networks where the nodes are mobile and end-to-end paths may not exist between any two nodes in the network. Each node is assigned to a particular group or domain administratively and will not change its group membership. A group may be a military platoon or a disaster rescue team. Each group of nodes may be confined to move only within a certain geographical area, e.g. ,different platoons may be sent to monitor different sub-areas of a bigger geographical area. The nodes within a group may either move independently within that assigned sub-area or move as a group within each assigned sub-area. When the nodes move as a group, each node will be located within a certain radius from its group center. This group center may be the location of its group leader or simply the geographical center of that assigned sub-area. We assume that each group has at least one node that is equipped with GPS and hence each group can receive location information through the Global Positioning System. We assume that each group has its unique group identifier which is included in the beacons messages that a node periodically broadcasts. For security reason, the beacon message may be encoded with different keys so that only those friendly groups can decipher a node's group identifier. Thus, a node can know if it hears nodes from other domains.

We assume that all groups may have one or more special nodes that are referred to as ferries. The ferries can move faster than each individual node or each group. These ferries are used to deliver messages between different groups. We assume that the different groups are friendly to one another and hence willing to deliver messages from any groups. We also assume that via long range radio communications, the ferries can be made aware of the locations of each group. Such an assumption is not unreasonable, e.g., military or rescue teams often have some long range radios that allow them to transmit or convey their group location information. We assume that such long range radios have limited bandwidth and hence are not used for regular message delivery.

3.2 Interdomain Routing Schemes

3.2.1 Gateway-Based Interdomain Routing (GBIR) Scheme

Figure 1 illustrates how GBIR works. There are three components in GBIR, namely (a) leader selection and transfer, (b) gateway registration, deregistration, and transfer, and (c) data delivery. All nodes (leader, gateway, regular) are assumed to support DTN stored-and-forward functionality.

We first state the assumptions we make for this scheme. We assume that the group can determine the location of its group's center ,e.g., a group leader may have a GPS that determines its location, and such information is then broadcasted to the whole group periodically. We also define the area within one hop from the boundary of the area which the nodes are allowed to move to be the boundary area. For example in Figure 1(d), we assume the nodes within each group can move only within the group's assigned quadrant. So, the shaded rectangular strip is considered the boundary area.

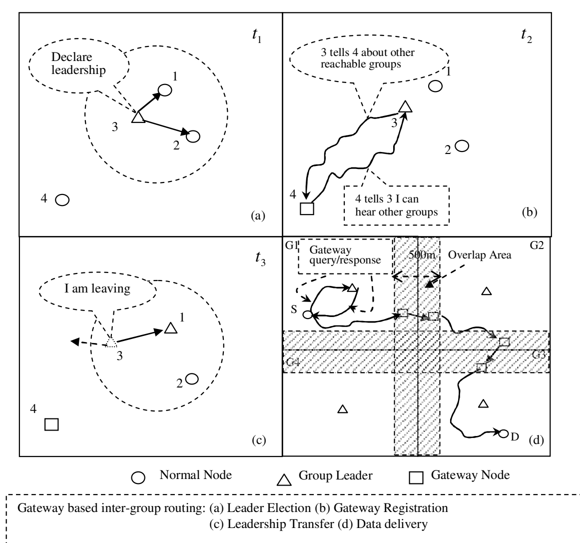

Fig. 1: Gateway based inter-group routing

(a)Leader Selection and Transfer

All nodes which are within one hop from the group's center will be leader candidates. At the bootstrap stage, every leader candidate backoffs a random time before sending a message to claim its leadership. The node that sends the claim first will be selected as the leader. If more than one leader succeeds almost the same time, the one closest to the group's center will be selected. The leader periodically sends out leader beacons. If the nodes located within one hop from the center do not hear a leader beacon for L consecutive periods, they assume the leader dies and can initiate a leader selection process again.

When a leader moves out of the one hop area from the center, it checks all its neighbors and chooses the one which is closest to the group's center to take over its leadership.

(b) Gateway registration, deregistration and transfer

Each node broadcasts a beacon periodically when it enters into the boundary area. Any node that hears messages from other groups forwards a gateway registration message to its group leader. Since its group leader is always within a single-hop from the group's center, geographical routing will be used to forward registration messages. Thus, even if the gateway does not know the identity of its current leader, it can still register successfully with its leader. The registration message contains gateway location information.

In Figure 2(a), we illustrate how two groups which move as a group discover one another. A node, n1, within Group 1 (with group center GC(1) starts broadcasting beacons when it enters into the shared area (also known as the boundary area). Nodes within a group are pre-configured with either the location of the group's boundary or a value k that says that they are in the boundary area if they hear a beacon from the group leader (at the group center) with a TTL more than k. At time t1, the two groups are far away so n1 and m3 cannot hear one another's beacons. At time t2, they can hear one another. So, node n1 (m3) becomes the gateway of Group 1 (Group 2).

When a registered gateway cannot hear from other groups, it sends a de-registration message to its leader. When a gateway node moves away from the boundary area, it finds a neighbor from its subnet that is currently in the boundary area to take over.

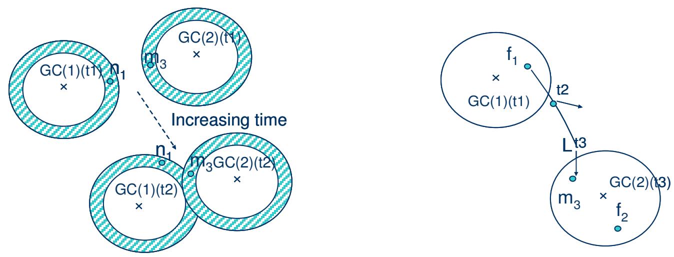

(a) GBIR (b) FBIR

Fig. 2: Discovering Other Groups

(c) Data Delivery

When a node has data to send to another group, it queries its leader for the gateway information. The leader only provides information on a rough location of the gateway. When a query node gets a response from the leader, it uses the underlying multihop routing or geographical forwarding scheme to deliver the data packets to the gateway.

3.2.2 Ferry-Based Interdomain Routing (FBIR) Scheme

In FBIR, we assume that each group has one or more ferries which are responsible for delivering inter-group messages from this particular group to all other groups. We further assume that each group member knows the identifier of its own group's ferry and use intradomain routing protocol to deliver inter-domain traffic to the ferry. We further assume that the ferries communicate among themselves via long range radios so they know the rough location of every group.

A ferry can be in either of two possible states: (i) local i.e. being with its own group, and (ii) roaming i.e. the ferry is visiting other groups. We also assume that the ferry is willing to deliver messages from other groups when it is in the roaming state. When a ferry crosses over to a neighboring territory, it broadcasts a service announcement message periodically to discover nodes from other groups. A ferry also periodically checks the packets that are stored in its buffer to see if (a) there are packets which have been queued for more than w seconds, or (b) the buffer occupation is full. If either of these conditions happens, then the ferry will start to move towards the destination group of the oldest message among those queued messages. If the second condition is the triggering condition, then the ferry will visit the destination group with the highest number of queued messages. If both conditions are triggered, then the ferry will consider condition (b) to be more important than condition (a) and act accordingly.

Once the ferry leaves its own group, it needs to discover a node from the destination group. Figure 2(b) illustrates how the ferry f1 from Group 1 discovers a node m3 from Group 2 (the destination group). The ferry f1 travels towards the location of Group 2 that it last heard. The ferry issues hello messages periodically to look for nodes from the destination group. When ferry f1 reaches the location L at time t3, it can hear the beacon from m3 and realizes that it has reached Group 2. Once such a node is discovered, the ferry transfers all messages destined to the destination group to this newly discovered node. This newly discovered node then delivers the messages to the destination nodes using its own group's intradomain routing scheme.

The visiting ferry will then leave immediately for another destination group if there are messages destined to other groups in the ferry's buffers that satisfy one of the triggering conditions. Otherwise, the ferry stays in the visiting group to collect messages destined to the ferry's own group. The ferry continues the data collection until one of the above two conditions triggers the ferry to move to the next group or move back to its own group.

Each node within a group delivers an intergroup message either to the local ferry or a destination group's ferry that is currently visiting. Figure 3 illustrates how the ferry-based interdomain routing scheme operates. Condition (a) triggers F1 to move to Group 3. After delivering the messages for G3, F1 performs a random walk to collect messages for G1 until the oldest message for G4 has been queued for w sec, and hence triggers F1 to move to G4 before eventually moving back to G1

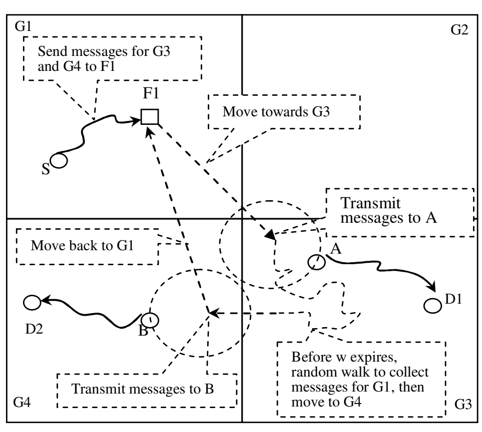

Fig. 3: Ferry-Based Interdomain Routing Scheme

4 PERFORMANCE EVALUATION

4.1 Simulation Setup

In order to compare the two interdomain routing schemes, we implement these two schemes in ns-2 [18]. We also implement two intradomain routing protocols, namely (a) a multihop routing with custody transfer (MRCT) scheme, designed for DTNs [19] and a two-hop routing protocol described in [13]. The MRCT scheme operates like DSR except that a route request message can be sent by either the source or an intermediate node to find a route to the destination. The other big difference between this MRCT scheme and DSR is that intermediate nodes which act as custodians will store the packets until they can find a next-hop custodian or destination to transfer the stored packets. A node will store a message in its buffers until it receives an acknowledgement from the next-hop custodian that that node has received the message. For the two-hop relay scheme, a packet is encoded by the source into n blocks and delivered to the first n contacts. These first n contacts will store the packets and deliver them only when the contacts can communicate with the destination node. The destination node only needs to receive k blocks for message reconstruction. In our simulation, we set (n=8, k=4) for the two–hop relay scheme.

Since we focus our study on the delivery performance of intergroup messages, the performance metrics used in our evaluation are:

- Delivery Ratio (DR), which is the successfully received number of intergroup messages divided by the total number of intergroup messages sent.
- Average Delivery Latency (DL) which is defined as the average end-to-end delay incurred by intergroup messages, and
- Transmission Efficiency (TE) is the total number of delivered intergroup messages (measured in terms of bytes) over the number of transmissions (which includes control overhead mes sages) used to deliver these intergroup messages. Each control message is assumed to be 35 bytes. Control messages are those that are sent to find routes e.g. route requests/response for intradomain routing or periodic beacons messages.

When comparing two schemes A and B, we say Scheme A performs better than Scheme B if $DR_A > DR_B$ or $DR_A \approx DR_B$, $TE_A > TE_B$ or ($DR_A \approx DR_B, TE_A \approx TE_B, DL_A \prec DL_B$) .

In our simulation, we use a few network scenarios (as shown in Figure 4): (a) the default network scenario (NS1) consists of four groups of nodes. Each group has 20 nodes which are constrained to move within a geographical area of 1400x1400m², (b) Network Scenario 2 (NS2) where the four group of nodes are allowed to roam the whole area of 2800x2800 m², (c) Network Scenario 3 (NS3) where different number of groups (each group has 10 nodes) are allowed to move within a square area. The size of the area is chosen based on the number of groups so as to keep the average node density to be the same as that in NS1 and NS2. For Network Scenario 2, the nodes within a group can either move independently or as a group as shown in Figure 4.

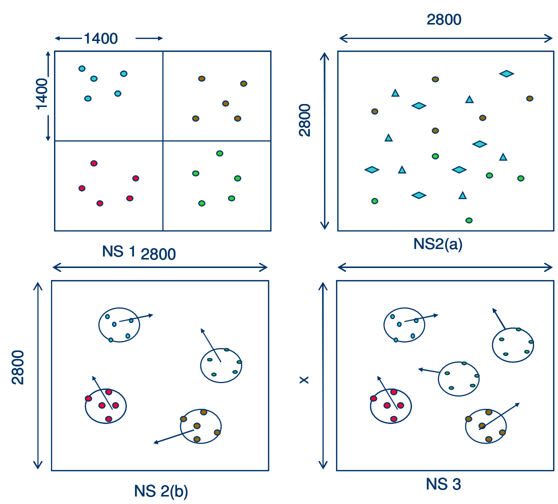

Fig. 4: Network Scenarios

4.1.1 Mobility Model

The default mobility model used is the random waypoint (RWP) model [17] but we also incorporate three additional mobility models, namely (i) Zebranet mobility model [13], (ii) Random Point Group Mobility (RPGM) model [20], and (iii) UMassBusNet mobility model [24]. With the random waypoint model, unless otherwise stated, we set the pause time to be 10 seconds, and the maximum node speed to be 5 m/s. For the Zebranet model, we create a semi-synthentic Zebranet mobility model as follows: we synthesize node speed and turn angle distributions from the observed data and create other node-movements using the same distribution. We use both distance and time scaling to fit the original data found in the trace into the network environment that we are interested in. For the UMassBusNet model, we extract the locations of twenty buses at different times in one trace, and scale their relative locations to fit into the 1400x1400m^2. This will be used to represent the node movement in one group. We repeat this process four times for four groups. The RPGM model is used to model military battlefield communications. Here, each group has a logical center (group leader) that determines the group's motion behavior. Initially, each member of the group is uniformly distributed in the neighborhood of the group leader. Subsequently, at each sampling instant, every node has a speed and direction that is derived by randomly deviating from that of the group leader. For the RPGM model, we use the mobility generator in [21] to generate the group mobility trace. The velocity of each member is characterized as follows:

(i) |Vmember(t)|=|Vleader(t)|+random()*SDR*max_speed,

(ii) θmember(t)=θleader(t)+random()*ADR*max_angle.

The speed deviation ratio (SDR), and the angle deviation ratio (ADR) are both chosen to be 0.1 in our simulations. All nodes within a group are located within a circle with a radius of 700 m centered at the group leader. The group leaders are allowed to move either within their own 1400x1400 m^2 area or within the whole 2800x2800 m^2 area.

4.1.2 Traffic Model

For inter-group flows, we use the random group communications where each group has one intergroup flow to each of the other three groups. We select a random node that belongs to the selected source or destination group to be the source or destination of each inter-group flow. Thus, there are 12 intergroup flows when there are four groups. Each flow generates CBR traffic with a packet size of 512 bytes. The traffic generation rate for each flow is varied from 0.2 msgs/sec to 2 msgs/sec. For scenarios where we include intradomain flows, the sources and destinations of such flows are randomly selected among the nodes from a group.

All the nodes communicate using a transmission range of 250m and a bandwidth of 2 Mbps. For FBIR, there is only one ferry in each group with its speed set to 15 m/s and w is set to 1000 seconds. The ferry and the regular node buffer sizes are set to 600 messages each. Messages are dropped according to Last In

First Out (LIFO) policy. We run each simulation for 10,000 seconds with a warming up period of 1000 seconds and the reported simulation results are based on the average of 5 runs.

4.2 The Usefulness of Interdomain Routing Schemes

In this section, we identify scenarios where using interdomain routing scheme will provide better intergroup delivery performance over the flat routing scheme. We first use NS1 where there are four groups of nodes. Each group has 20 nodes which are distributed over its own 1400x1400 m^2. We use random group communication patterns with each group having a flow to the other 3 groups. The routing schemes for delivering messages for these 12 intergroup flows are (a) multihop routing with custody transfer (MRCT) [19], (b) GBIR, and (c) the FBIR schemes. We assume that when a flat routing scheme like MRCT is used, all 80 nodes belong to the same group. For hierarchical routing approach, MRCT is used as the intradomain routing scheme while either FBIR or GIBR is used as the interdomain routing scheme. We vary the message rate in our experiments. Figures 5(a)-5(c) plots the delivery ratio, the average delay, and the transmission efficiency.

From the results, we see that both of the interdomain routing schemes achieve higher delivery ratio, lower average delay and higher transmission efficiency compared to those achieved using the flat routing scheme. Such results are not surprising since the flat routing scheme not only takes longer to discover an interdomain route but also uses more hops to deliver intergroup packets. Routes with more hops usually break more often and extra overhead is incurred for route re-discovery. Thus, the flat routing scheme has poorer delivery performance.

FBIR and GBIR achieve comparable delivery ratio (see Fig 5(a)) but FBIR achieves lower average delay than GBIR (see Fig 5(b)). FBIR also achieves higher transmission efficiency than GBIR. This is understandable since FBIR uses fewer hops to deliver intergroup traffic, and the ferry travels at five times the average node speed. On the average, FBIR takes a total of 13 hops (6 hops from a source node to the ferry, one ferry hop, and another 6 hops from the recipient node to the destination) to deliver intergroup messages but GBIR takes a total of 19 hops (9 hops in each domain plus a gateway-gateway hop) to do so.

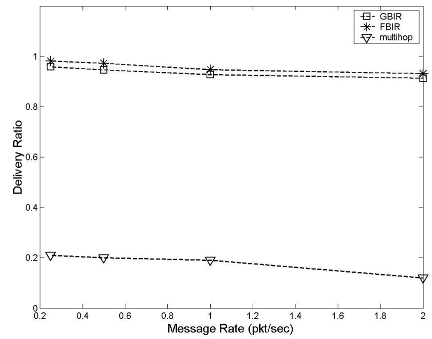

Fig. 5(a): Delivery Ratio vs Message Rate (NS1)

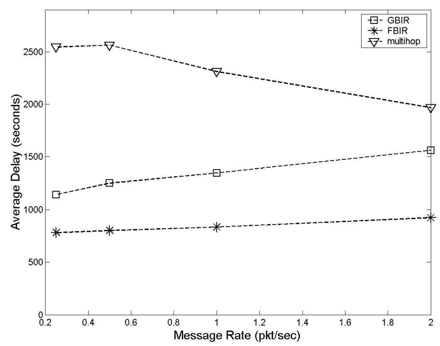

Fig. 5(b): Average Delay vs Message Rate (NS1)

From Figure 5(b), we see that the average delay for the flat routing scheme drops with increasing message rate. This is misleading and can be explained as follows: the messages that require more hops to be delivered are more likely to be dropped due to buffer overflow. The successfully delivered messages as message rate increases are usually those that incur fewer hops, and hence the average delay drops. We also observe that the average delay for both FBIR and GBIR increases with increasing message rate but the rate of increase is higher for GBIR. This can be explained as follows: the increasing E2E delay is due to the increasing access delay at each hop, and since GBIR takes more hops to deliver intergroup messages, its rate of increase is higher. The transmission efficiency for both FBIR and GBIR increases slightly with increasing message rate because more messages can be delivered during the contact duration before delivery path changes.

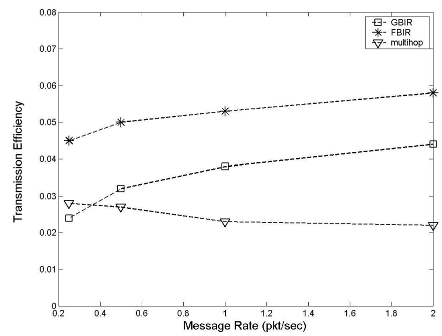

Fig. 5(c): Transmission vs Message Rate (NS1)

In the above experiment, the four groups of nodes are restricted to move within a pre-assigned area and hence it is less likely for the nodes from one group to meet nodes from another group. We are interested in understanding whether there is any scenario where flat routing will perform better. We suspect that flat routing may perform better only when nodes from different groups are mixed together. Thus, we repeat the above experiment using network scenario NS2(a) where nodes are allowed to move independently within 2800x2800m^2.

In Figures 6(a) to 6(c), we plot the results we obtain. We see that flat routing performs the best followed by GBIR and FBIR. The flat routing achieves the best delivery ratio followed by GBIR and FBIR. The paired t-test results [22] indicate that the difference in delivery ratios between GBIR and FBIR is not significant at low load but very significant (P value of less than 0.0001) at higher load (> 1pkt/s). In this NS2(a) scenario, flat routing scheme delivers messages with the smallest hop counts. However, both GBIR and FBIR take longer delivery paths. In this scenario, the average numbers of hops taken using flat routing, GBIR and FBIR are 3.5, 6.4 and 11.5 respectively. This is not surprising since an intergroup message using FBIR needs to be delivered first to the source group ferry, then to a node in the destination group before it will be delivered to the destination node. Similarly, an intergroup message using GBIR needs to be delivered first to the source group gateway, then to the destination group gateway before being delivered to the destination node.

The above results reveal that interdomain routing scheme is useful when different groups are segregated from one another. To verify further this conclusion, we repeat our experiment using the network scenario NS2(b) where each group of nodes moves together using the RPGM model. We plot the results in Figures 7(a) to 7(c). We see that as in the NS1 scenario, FBIR performs the best, followed by GBIR and flat routing when network scenario NS2(b) is used.

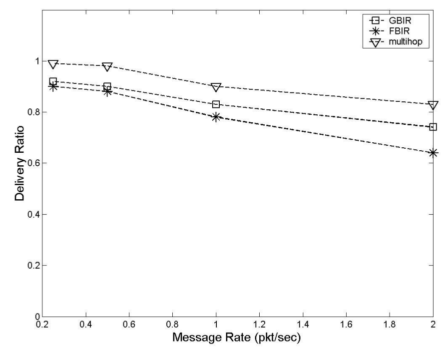

Fig. 6(a): Delivery Ratio vs Message Rate (NS2(a))

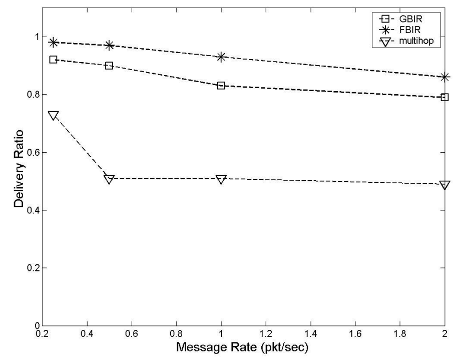

Fig. 7(a): Delivery Ratio vs Message Rate (NS2(b))

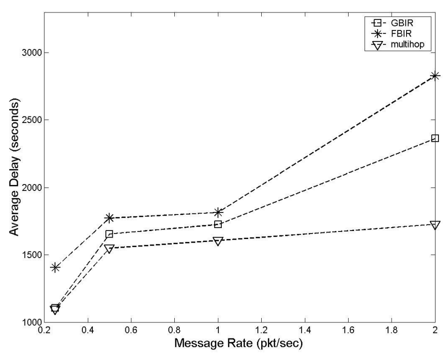

Fig. 6(b): Average Delay vs Message Rate (NS2(a))

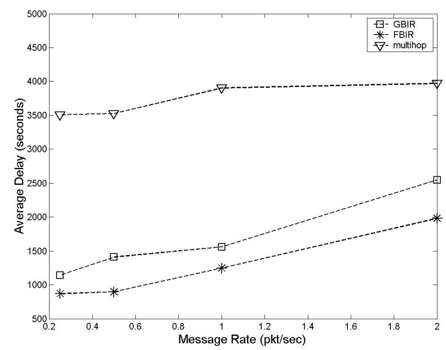

Fig. 7(b): Average Delay vs Message Rate (NS2(b))

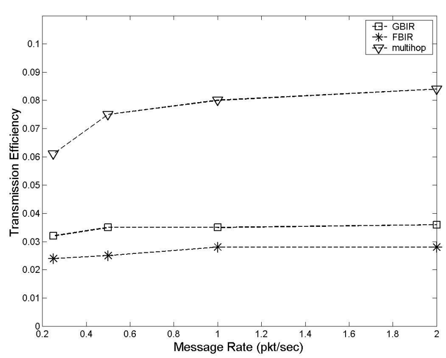

Fig. 6(c): Transmission Efficiency vs Message Rate (NS2(a))

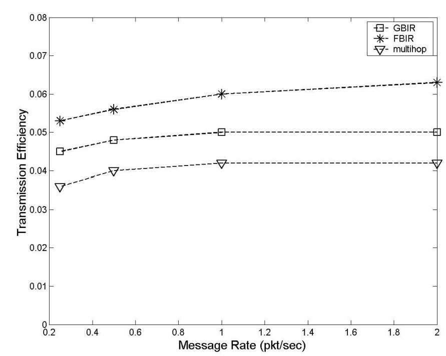

Fig. 7(c): Transmission Efficiency vs Message Rate (NS2(b))

In addition, we make several interesting observations when we compare the delivery performance in NS1 and NS2(b). Compared to the NS1 scenario, flat routing performs better in NS2(b) since nodes are more likely to discover other groups when their group movement is not constrained to a certain geographical area. Another observation is that the delivery performance for both FBIR and GBIR are poorer in NS2(b) (when compared to NS1). In NS2(b), each group is allowed to move within a bigger area ($2800 \times 2800 m^2$). When FBIR is used, the ferry will take longer time to complete a visiting trip in NS2(b). Thus, more messages will be queued at the source nodes while the ferry is away which means more messages may be dropped in NS2(b) than in NS1. Similar argument can be made for GBIR. It takes longer for a gateway to find a gateway from another group in NS2(b), and the intergroup message delivery path also takes more hops. Both these factors contribute to more message queuing which result in more buffer overflows. We will discuss more on the impact of mobility models on FBIR in a later section.

4.3 Impact of different intradomain routing schemes

Using the random group communication traffic pattern, we investigate the impact of different intradomain routing schemes on the performance of the interdomain traffic delivery. We repeat the experiment using network scenario NS1, but letting all four groups use two-hop relay scheme as their intradomain routing schemes. Figures 8(a) to 8(c) plot the delivery ratio, the average delay and the transmission efficiency with all four groups use either the MCRT or the two-hop relay scheme as the intradomain routing scheme.

From the plots, we see that the interdomain traffic delivery performance is better when all groups use the MRCT scheme as the intradomain routing scheme. With FBIR as the interdomain routing scheme, the intergroup message delivery ratio drops from 97%to 83% (see Figure 8(a)) as the message rate increases from 0.25 pkt/s to 2 pkt/s when two-hop relay scheme is used for intradomain routing. However, if the MRCT scheme is used, the delivery ratio only drops from 98% to 95% for the same increase in message rate. The explanation is as follows: two-hop relay scheme takes longer time to find the ferry and hence need to queue messages longer. Hence, more messages will be dropped when two-hop relay scheme is used as the intradomain routing scheme. A similar trend is observed when GBIR is used as the interdomain routing scheme.

For average delay, we see from Figure 8(b) that using two hop relay scheme as the intradomain routing scheme causes higher E2E delay increase when GBIR is used as the interdomain routing scheme. With

FBIR as the interdomain routing scheme, the E2E delay increases by 15-80% when the two hop relay scheme rather than the MRCT scheme is used as the intradomain routing scheme. With GBIR as the interdomain routing scheme, the E2E delay increases by 86 to 97% when the two hop relay scheme rather than the MRCT scheme is used as the intradomain routing scheme. The larger delay increase observed with the GBIR scheme can be explained as follows: it takes longer time for a node to find the gateway (and vice versa), and for the two gateways to find one another using the two-hop relay scheme.

As for transmission efficiency, we see from Figure 8(c) that the two hop approach achieves higher transmission efficiency since it takes fewer hops to deliver the message. Using FBIR as the interdomain routing scheme, a total of 5 hops is needed to deliver an intergroup message (2 hops from the source to the ferry, one ferry hop and another 2 hops at the destination group) when two hop relay scheme is used as intradomain routing scheme while 13 hops are needed (6 hops from the source to the ferry, one ferry hop and 6 hops at the destination group) when multihop scheme is used. Again, a similar trend is observed when GBIR is used as the interdomain routing scheme.

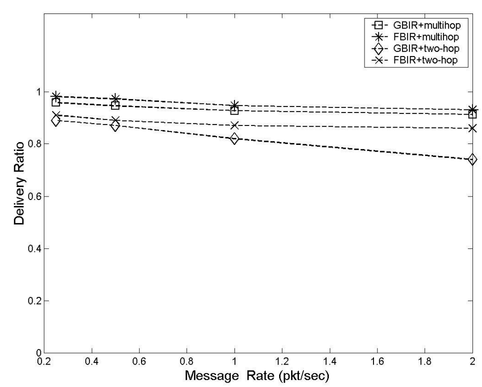

Figure 8(a): Delivery Ratio vs Message Rate (NS1)

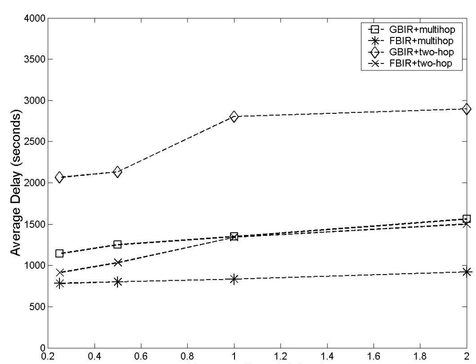

Figure 8(b): Average Delay vs Message Rate (NS1)

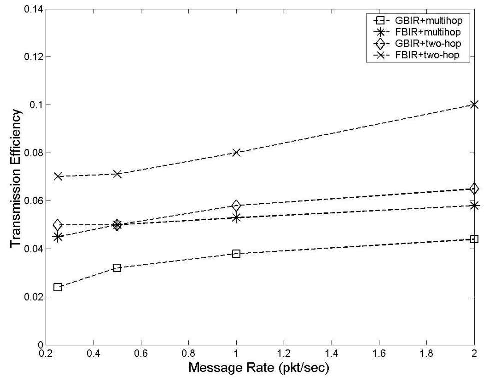

Figure 8(c): Transmission Efficiency vs Message Rate (NS1)

4.4 Impact of Ferry Speed

In scenarios where different groups can potentially hear one another, we suspect that FBIR performs better than GBIR because the ferry travels faster than the average node speed. Thus, in this section, we conduct an experiment using the network scenario NS2(b) but vary the ferry speed from 3 m/s to 20 m/s and compare the FBIR performance with GBIR. Figures 9(a) to 9(c) plot the delivery ratio, the average end-to-end delivery latency, and the transmission efficiency of the interdomain traffic as the ferry speed changes from 3 m/s to 20 m/s at two different traffic load (each flow generating 0.5 pkt/s or 1 pkt/s). We also plot the results obtained using GBIR.

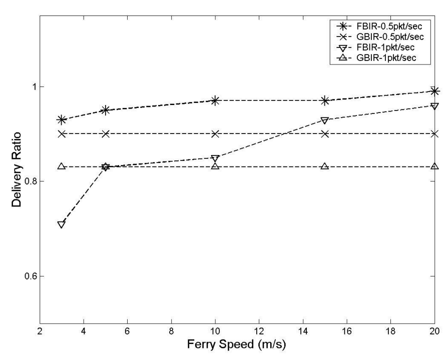

Figure 9(a): Delivery Ratio vs Ferry Speed (NS2(b))

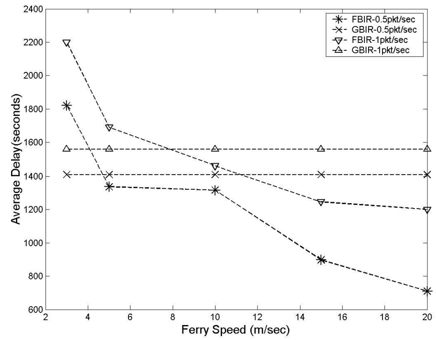

Figure 9(b): Average Delay vs Ferry Speed (NS2(b))

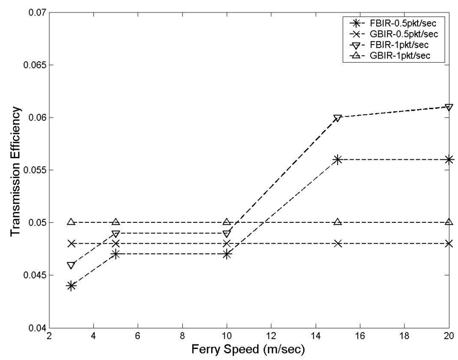

Figure 9(c): Transmission Efficiency vs Ferry Speed (NS2(b))

From Figure 9, we see that at low load (0.5 msg/s), GBIR performance is comparable to FBIR when the ferry speed is the same as the average node speed (3m/s). When the ferry speed is 3 m/s, the round trip time for ferry is about 970 sec. This means that a total of about 1455 new messages (970*0.5 msg/s*3 (flows)) will be generated while the ferry is away. If we assume that the messages can be stored at 3 source nodes, and the ferry, we have a total of 2400 buffers. Then, we will not see much buffer overflow using FBIR when each flow generates 0.5 msg/s. But when the traffic load is 1 msg/s per flow, then a total of about 2910 messages are generated while the ferry is away which means that 510 messages are likely to be dropped which result in a delivery ratio of about (510/(2910-510))=0.787 (simulation result is 0.71). On the other hand, GBIR does not suffer from this problem since messages can be queued at more intermediate nodes. Thus, we see that GBIR achieves higher delivery ratio than FBIR when the ferry speed is 3 m/s. The average delay using FBIR is also higher than GBIR when the ferry speed is 3m/s since messages are relayed faster via multiple hops in GBIR than being carried by the slow moving ferry in FBIR.

Figure 9(b) shows that when the ferry speed is 8m/s, the FBIR scheme will achieve similar average delay performance as in the GBIR scheme for the 1pkt/s scenario. This can be explained as follows: we observe that the average intradomain delay in GBIR is about 430s when the source rate is 1 pkt/sec. The mean intergroup meeting time can be derived using the Eqn (6) in [23] by setting the effective transmission range to be 825 m (700 is the group radius and 125m is half of the radio transmission range) to be 700s. Thus, the average end-to-end delay for GBIR with 1 pkt/s is about 1560s(=2*average intradomain delay+average intergroup meeting time). This mean delay value is also observed from our simulation results. For FBIR, the observed average intradomain delay is about 530s and the mean distance between the source and destination groups is about 1700m. When the ferry speed is 8m/s, then the average delay for FBIR will be 1485s(=2*avg intradomain delay+avg ferry round trip time =2*530+3400/8). Thus, as long as the ferry speed exceeds 8m/s, FBIR will perform better than GBIR even at a high load of 1 msg/s per flow. It will be useful to derive analytical expressions for the average intradomain delay and the average distance between two groups when the groups move according to the RPGM. Then, given any average group moving speed, one can determine the ferry speed that allows FBIR to perform better than GBIR. We leave this for future work.

4.5 Impact of Mobility Models

In this section, we investigate how the mobility models impact the delivery performance of the interdomain traffic. We use the random group communication traffic pattern but let the nodes move according to (a) RWP, (b) Zebranet (c) UMassBusNet, or (d) the RPGM mobility model. In the first three mobility models, the different groups are constrained to move within its own 1400x1400m^2 area but in (d) the different groups are free to move within the 2800x2800m^2 area. All four groups use the MRCT scheme as the intradomain routing scheme. Figures 10(a) to 10(c) show the results for delivery ratio, average delay and transmission efficiency obtained with the four mobility models. To help explain the results in Figures 10(a) to 10 (c), we provide CDF plots for contact duration and intercontact time for RWP, Zebranet, and UMassBusNet models in Figures 11(a) and 11(b) respectively. The average contact durations for RWP, Zebranet and UMassBusNet are 98, 41 and 243 seconds respectively. The mean intercontact times for RWP, Zebranet and UMassBusNet are 483, 405, 2342 respectively.

From Figure 10(a), we see that the two mobility models (RWP and Zebranet) achieve similar delivery ratio, RPGM model gives slightly poorer delivery ratio but UMassBusNet model achieves the poorest delivery ratio. This is not surprising since the mean intercontact times for UMassBusNet model is more than five times that of RWP and Zebranet. This means messages are queued longer when UMassBusNet model is used and thus buffer overflows happen more frequently. The RPGM model gives slightly poorer delivery ratio because it takes longer for the ferry to make a visiting trip when the different numbers of groups are allowed to move within a larger area (2800x2800m^2 vs the 1400x1400m^2 for the RWP and Zebranet models). Thus, messages are queued at souce/intermediate nodes longer and may cause new packets to be dropped when buffers overflow.

The average delay is smallest with RWP, followed by the Zebranet model, the RPGM model, and the UMassBusNet model. Again the large delay for UMassBusNet model is not surprising because of its high intercontact time. The delay observed in RPGM model is higher than RWP and Zebranet because it takes longer for the ferry to make a visiting trip when the different groups are allowed to move over 2800x2800m^2. We do not include confidence intervals of our simulation results in the plots so as not to clutter the curves. Our paired t-test results did show that the difference in the mean values between the RPGM/RWP pair, UMassBusNet/RWP pair or RPGM/Zebranet pair, UMassBusNet/Zebranet pair is very significant (P values less than 0.0001). The transmission efficiency for the UMassBusNet is higher because fewer hops are used to deliver messages in that scenario. The average number of hops observed for the RWP model is 13 but it is only 6 for UMassBusNet model.

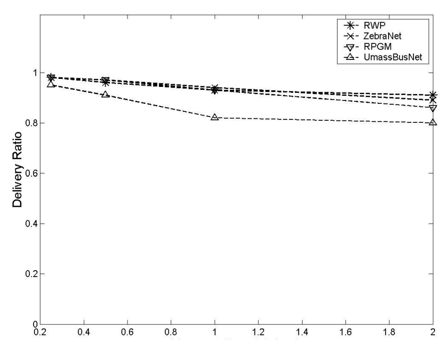

Figure 10(a): Delivery Ratio vs Message Rate (all 4 mobility models)

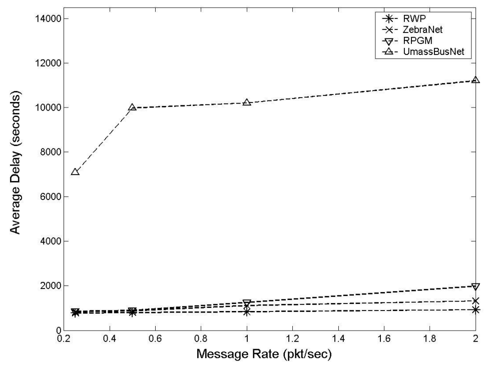

Figure 10(b): Average Delay vs Message Rate (all 4 mobility models)

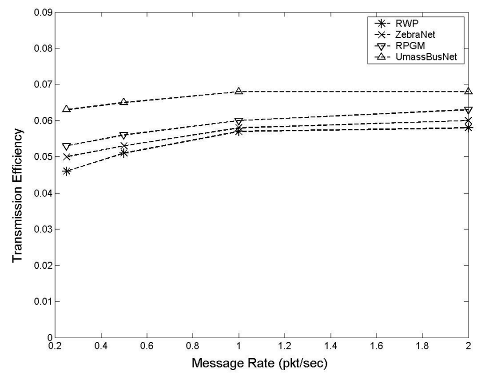

Figure 10(c): Transmission Efficiency vs Message Rate (all 4 mobility models)

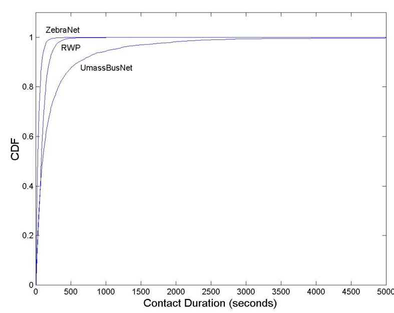

Figure 11(a): CDF of contact duration for RWP, Zebranet, UMassBusNet

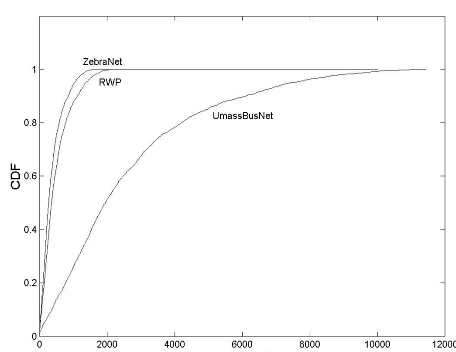

Figure 11(b): CDF of intercontact time for RWP, Zebranet, UMassBusNet

4.6 Impact of different number of groups

In this section, we investigate the impact of having different number of groups on the delivery performance of FBIR. For this set of experiment, we use network scenario NS3 where each group has 10 nodes. The node density is kept similar to the node density we used in earlier experiment as we vary the number of groups. With 4 groups, 6 groups, 8 groups and 10 groups, the total area covered by the groups are 1980x1980, 2425x2425, 2800x2800 and 3130x3130 m^2 respectively. Each group moves according to the RPGM model. All nodes within a group are located within a circle with a radius of 495m centered around its group leader. Each group has three intergroup flows to three randomly selected destination groups. Each intergroup flow generates a message rate of 1 msg/sec. We tabulate our results for the delivery ratio, the average end-to-end delay the transmission efficiency, average intradomain delay, average ferry round trip time and average number of end-to-end (E2E) hops in Table 2.

# of Groups	Delivery Ratio	Avg E2E Delay	Transmission Efficiency	Avg Intra-domain Delay	Avg Ferry Round Trip Time	Avg E2E Hops
4	0.92	1166	0.060	445	173	5.7
6	0.89	1203	0.059	451	215	6.1
8	0.86	1323	0.058	468	273	5.7
10	0.82	1413	0.059	481	316	6.1

Table 2: Results with different number of groups.

Our results indicate that as the number of groups increases, the delivery ratio decreases and the average delay increases. This can be explained as follows: as the geographical area increases, the ferry takes longer to complete its visiting trip and hence messages need to be stored longer at the source and intermediate nodes. This causes more buffer overflows and larger end-to-end delays. The transmission efficiency remains relatively constant with different number of groups. Our results indicate that FBIR is scalable with increasing number of groups. Of course, one can always increase the number of ferries per group to improve the delivery performance.

4.7 Impact of Buffer Size and Intradomain Flows

Next, we investigate how the buffer size affects the intergroup message delivery, we pick the scenario where there are 10 groups (since its delivery ratio is 82% with a per-node buffer size of 600 messages for regular nodes or ferry. Table 3 tabulates our results as the buffer size changes from 600 to 1000 messages. The results indicate that with increasing buffers, the delivery ratio improves but at the expense of increasing average end-to-end delay.

Buffer Size	DR	Avg E2E delay	TE	Avg Intra-domain Delay	Avg Ferry Round Trip Time	Avg E2E Hops
600	0.82	1413	0.06	481	316	6.1
800	0.86	1461	0.06	501	323	6.2
1000	0.93	1589	0.06	562	331	6.2

Table 3: Impact of Buffer Size (10 groups scenario)

We also investigate how the presence of intra-domain messages affects the intergroup message delivery performance. Table 4 tabulates our results for the 4-group scenario with 0, 2 intragroup flows/group, and 4 intragroup flows/group. The E2E delay is averaged over all delivered intergroup messages but the average intradomain delay includes the intradomain messages. The 4-group subscenario is chosen purely because it takes too long to simulate a scenario with large number of groups and high packet rate using the NS-2 simulator. The general trend is expected to be the same. Our results indicate that the presence of intradomain flows increases the end-to-end delay slightly due to increasing queuing delay.

Table 4: Impact of intradomain flows (4 groups scenario)

Intra-domain Flows	DR	Avg E2E Delay	TE	Avg Intra-domain Delay	Avg Ferry Round Trip Time	Avg E2E Hops
0	0.92	1166	0.06	445	173	5.7
2	0.92	1221	0.06	469	174	5.7
4	0.90	1273	0.06	492	179	5.8

5 CONCLUSION

In this paper, we describe two interdomain routing schemes for disruption tolerant networks, namely the gateway-based, and the ferry-based interdomain routing scheme. We have identified scenarios where using interdomain routing scheme to deliver intergroup messages will be useful. Our simulation studies show that both interdomain routing schemes provide better delivery performance when different groups are segregated from one another. In addition, we also show that the ferry-based scheme achieves higher delivery ratio and lower average end-to-end delay for the interdomain traffic than GBIR when the ferry speed is higher than the average node speed. It will be useful to derive some analytical expressions that allow us to determine the ferry speed which will make FBIR perform better than GBIR when the different groups move according to RPGM. We leave this for future work.

We have also demonstrated that the choice of intradomain routing scheme affects the delivery performance of intergroup messages. Using a multihop intradomain routing scheme is better than using a two-hop relay routing scheme. In addition, we investigate how mobility models impact the intergroup delivery performance of the FBIR scheme. We observe that the UMassBusNet model results in the poorest delivery performance followed by the RPGM, the Zebranet and the RWP models. The poorer performance of the UMassBusNet mdoel is due to the larger intercontact times, and the more frequent network partitions among different groups. We also investigate the scalability of the FBIR scheme. Our results with different number of groups indicate that FBIR scales well and can achieve 82% delivery ratio even with 10 groups of nodes. We also show that increasing buffer sizes at the ferry and regular nodes can improve the delivery performance.

6 ACKNOWLEDGEMENTS

This work has been supported by DARPA under Contract W15P7T-06-C-P430. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the sponsor of this work. We wish to thank Dr. Mark-Oliver Stehr for some useful comments on an earlier version of this paper.

7 REFERENCES

[1] D. Johnson, D. Maltz, “Dynamic Source Routing in Ad-Hoc Wireless Neteworks”, Proceedings of ACM Sigcomm, August 1996

[2] C. Perkins, E. Royer, “Ad-hoc On-Demand Distance Vector Routing”, Proceedgins of ACM, WoWMoM, Seattle, WA, pp 26-33, August 1999.

[3] S. Das, C. Perkins, E. Royer, “Performance Comparision of Two On-Demand Routing Protocols“, Proceedings of IEEE Infocom, March, 2000.

[4] K. Xu, M. Gerla, “A Heterogeneous Routing Protocol Based on a new stable clustering scheme”, Proceedings of IEEE Milcom, 2002.

[5] W. Ma, M. Chuah, “Comparisons of Interdomain Routing Schemes for heterogeneous ad hoc networks », Proceedings of ACM WoWMoM, 2005.

[6] Forrest Warthman, “Delay-Tolerant Networks (DTNs)”, DTN Research Group Internet Draft, March 2003.

[7] K. Fall, “A delay-tolerant network architecture for challenged internets”, in SIGCOMM, 2003.

[8] K. A. Harras, K. C. Almeroth, ”Inter-Regional Messenger Scheduling in Delay Tolerant Mobile Networks”, Proceedings of IEEE WoWMoM, 2006.

[9] M.M.B.Tariq, M. Ammar, and E. Zegura, “Message Ferry Route Design for Sparse Ad hoc Networks with Mobile Nodes”, ACM MobiHoc , May22-25, 2006.

[10] S.Jain, K.Fall, and R. Patra, “Routing in a Delay Tolerant Network”, SIGCOMM’04, Aug. 30-Sept. 3, 2004.

[11] J. Burgess, B. Gallagher, D. Jensen, and B.L.Levine, Maxprop: Routing for vehicle-based disruption-tolerant networks. In INFOCOM, 2006.

[12] A. Lindgren, A.Doria, and O.Scheln, Probabilistic Routing in Intermittently Connected Networks. In Proc. Workshop on Service Assurance with Partial and Intermittent Resources, August 2004.

[13] Y. Wang, S. Jan, M. Martonosi, and K. Fall, "Erasure-Coding Based Routing for Opportunistic Networks", Proceedings of ACM Sigcomm WDTN Workshop, August 2005.

[14] S. Jain, M. Demmer, R. Patra, K. Fall, "Using Redundancy to cope with Failures in a Delay Tolerant Network", Proceedings of ACM Sigcomm, August, 2005.

[15] Y. Sun, E. Royer, C. Perkins, "Internet Connectivity for Ad hoc Mobile Networks", International Journal of Wireless Information Networks, Special Issue on Mobile Ad hoc Networks, Vol 9(2), April, 2002.

[16] M. Ergen, A. Puri, "MEWLANA – Mobile IP Enhanced Wireless Local Area Network Architecture", Proceedings of IEEE VTC, September, 2002.

[17] T. Camp, J. Boleng, and V. Davies, "A survey of mobility models for ad hoc network research", Wireless Communications & Mobile Computing (WCMC), Vol .2, no. 5, pp. 483-502, 2002.

[18] "The network simulator ns-2", [Online] at http://www.isi.edu/nsnam/ns/.

[19] M. Chuah, P. Yang, B. Davison, L. Cheng, "Store and Forward Performance in a DTN", poster, Proceedings of VTC Spring, May, 2005.

[20] X.Hong, M. Gerla, G. Pei, and C.C.Chiang, "A group mobility model for ad hoc wireless networks", proceedings of ACM International Workshop on Modeling, Analysis, and Simulation of Wireless and Mobile Systems (MSWiM), August, 1999.

[21] F. Bai, N. Sadagopan, A. Helmy, "The Important Framework for analyzing the impact of mobility on performance of routing for Ad hoc Networks", Ad Hoc Networks Journal, Elsevier Science, Vol 1, Issue 4, pp 383-403, Nov, 2003.

[22] W. Navidi, "Statistics for Engineers and Scientists", McGraw Hill, 2006.

[23] T. Spyropoulos et al, "Performance Analysis of Mobility-assisted Routing", Proceedings of ACM Mobihoc, 2006.

[24] X. Zhang, J. Kurose, B. Levine, D. Towsley,H. Zhang,"Modeling of a bus-based disruption tolerant network trace", to appear in Proceedings of ACM Mobihoc, 2007.